403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:1

Appendix A

Intrusion Detection in the DMZ

Solutions in this chapter:

- **Intrusion Detection 101**
- **Repelling the Hacker**
- **Cisco IPS**
- **Snort**
- **The Poor Man's IDS**
- **More IDS Deployment Strategies**
- **Lessons Learned**

- ☑ **Summary**
- ☑ **Solutions Fast Track**
- ☑ **Frequently Asked Questions**

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:2

Introduction

In order to defend against attacks, we must first define just what the threats are. A *threat* is an intentional or unintentional act against something of value. Put simply, a threat is something bad that might happen to the item of value. The item of value could be a router, a firewall, or data itself. A *vulnerability* is a point of weakness in the defenses of the items of value. An attack (whether a passive or an active attack) is the act of finding or acting on the vulnerabilities of a system.

Today the firewall has become something of a commodity and is sold virtually everywhere. The marketing of the firewall has caused an indirect problem: Many people who purchase this commodity think that by using the firewall out of the box, the job of securing their network is completed. The reality is that their job has just begun, and if the task of securing the network is not finished, the entire network is still at risk of compromise. In today's world of threats from worms, viruses, intentional attacks, industrial espionage, and worse, a firewall is not enough.

Still, people think, "Yes, my network is protected—I installed a firewall." The disturbing truth is that configuring the firewall is not very easy yet, even with the various Web and GUI front ends. You still have to have some knowledge of how packets travel across a network and what makes a good packet turn into a bad packet. In order to protect your network effectively, you need to know the basics of how IP should work and how IP can be twisted to nefarious purposes, such as the SYN and FIN scans we saw in Chapter 15.

With a misclick of the mouse or a press of the wrong button, your shiny new firewall could suddenly have a very large hole in it due to misconfiguration. Although it's not technically a hole, this term is the name we append to a vulnerability; it's a "hole" in the system. Hackers live for misconfigurations, because they create the easiest targets to find and exploit. The biggest problem with the misconfigured firewall is the false sense of security it provides while you are sitting there thinking life is good and your personal or business assets are secure. Even with a properly configured firewall, an attacker can still get past it in many ways. Between HTTP attacks, Unicode attacks, buffer overflows, SQL injection, Trojan horses, viruses, worms, cross-site scripting, ftp bounce attacks, fragmented packet attacks, peer-to-peer tools and the like, there are many ways to get past a firewall.

To help understand how to effectively protect the network, we need to view the network in a different way. With the firewall at the border, the network is much like an egg, with a hard shell on the outside but an inside that is soft, unprotected, and easily affected if attacked. Think about what happens to the egg when "attacked" (as when you crack it on the lip of a bowl): It breaks, and everything inside comes running out. By the time your supposedly secure data is leaking past the firewall, it is much too late. We want our network to be more like an onion, with layers of defenses down to the last ring—which explains our defense-in-depth conversation from the last chapter. To this end, we use firewalls, IDS, honeypots, VLANs, ACLs, and applications such as antivirus programs or spyware detection programs like Spybot all together to provide a deep security posture. This posture allows us to

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:3

not only secure our networks and systems from many different threats, both internal and external, but to provide an "umbrella" of security methods to create the "onion." Along with the technology, we use a layered architecture of network design. We have the perimeter network defenses, then we move to the DMZ, incorporating the various resources that need to touch the Internet, such as Web, e-mail, and FTP servers. Behind those resources are more routers or firewalls and then the internal or protected network. But even on the protected network we use VLANs, more routers with ACLs, and physical security. In a well-defended network, the attackers should not get past the DMZ, which minimizes any damage the attacker can cause.

A firewall is the drawbridge of the castle. On one side of the drawbridge are the hordes of barbarians, and on the other side is the castle. You can raise or lower the drawbridge based on need or in the case of a firewall, you can open or close ports. You can block packets based on source or destination IP address and by network protocol. You can hide IP addresses by using Network Address Translation (NAT) or filter frames based on Media Access Control (MAC) addresses. The newer firewalls sold today will act as proxies between IP addresses and offer a higher level of network protection than the firewalls of just a few years ago. Another thought here is that as time passes, the single role of the firewall slowly disappears, and we start to see many different solutions offered in a single product. For instance, most firewalls today are very modular and will perform firewall, VPN, and IDS services, all in one. SnapGear manufactures an appliance aimed at small businesses that does just this thing. With all this in mind, you should also be aware that firewalls do have their limits. Although the firewall can block a protocol, what do you do about an attack that is hidden inside of a permitted protocol such as HTTP or FTP or SMTP or DNS? Enter the IDS.

Intrusion detection systems are a relatively new technology in the corporate and personal computing security realm. They arrived on the scene in the mid-1980s or so and were mostly research projects or military projects. With the Internet becoming a critical part of daily life, IDS has taken a more prominent role in networks over the past 5-6 years. Now IDS is a constantly evolving technology, sometimes ahead of the attack and sometimes behind it. There are commercial versions of IDS, such as the older Cisco NetRanger, which is now the Cisco IPS sensor, and open source IDS systems such as Snort, which can be found at www.snort.org. Along with the IDS come signature files and various graphical front ends, in addition to various databases for storing log files.

An IDS sensor works in simple terms, much like an alarm system for a house. It watches to see if anything trips the alarm and if so, logs the event and sends the alarm to the systems administrator. The IDS watches the network traffic as it goes by and looks to see if it can match a signature or detect a protocol or behavioral anomaly. The IDS is a much more elegant solution to the problem of hidden attacks than the deny-all or accept-all management of packets or protocols that the traditional firewall solution offers. However, the IDS does not work by itself - it supplements the firewall's protection.

There are many obstacles to deploying an effective IDS solution. Problems range from report generation to where to place the sensor(s) or how to diminish the false positives that occur, while preventing false negatives! The whole concept of the IDS is so important to the

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:4

enterprise DMZ because, unlike the firewall, which is a relatively fixed defense, the IDS can be an active defense. The IDS can look for and find problems or attacks within a permitted protocol such as HTTP. Many firewalls cannot isolate the payload of the HTTP traffic; the firewall can only tell if the traffic is port 80 and if it is permitted based on parameters such as IP address, direction, or if the packet is a fragment, and some can go further by distinguishing idiosyncrasies concerning the traffic stream, such as whether the packet is a FIN without the first part of the TCP handshake.

A new tool in the defense of your network's DMZ is the *honeypot*. The honeypot is the modern equivalent of the canaries that were used in mines to warn of unseen dangerous gases. In modern networks, the honeypot "sees" the unseen dangers of someone testing your network or actually attacking your network and alerts you—not by falling over dead like the canary but by alerts sent via technology such as syslog, cell phones, e-mail and pagers.

A honeypot is nothing more than a server set up in the DMZ to act, look, and feel like a legitimate target. However, the honeypot can be a nasty surprise for the attacker, with behavior ranging from acting as the early warning system of the network to a sophisticated attack logger. The Honeynet Project at http://www.honeynet.org refers to these two types of honeypot behaviors as "low interaction" and "high interaction". The honeypot can give out faked replies while logging keystrokes and then alerting and/or archiving files for later analysis. What better way to know if someone is sneaking around your house than to have them announced with a ringing doorbell? In this chapter, we examine the intricacies of using a honeypot in the DMZ so that we are not only warned about an attack after the fact but we learn how to supply misinformation and even log the hacker's activities. These log files can play a very important part in deciphering how the attack took place and then building a defense against it. The logs, when properly handled and preserved, can also provide legal forensic evidence about attempted attacks and provide a timeline of when the events took place.

A forgotten item in many security checklists is to look at the Web site content itself. Often programs and programmers insert company information and contact information into corporate Web pages. Along with phone numbers, names, and other personal information, you can often find directory information or even key internal IP addresses if the link has been hardcoded by a lazy programmer. All this information can be of help to a would-be attacker. This information is very easy to cull from the Web site by "ripping," or stealing, the entire site to a local drive using any of a number of common utilities, such as Teleport Pro, which can found at www.tenmax.com/teleport/pro/home.htm, or HTTrack Website Copier, found at www.httrack.com. Once the site is on the local drive, the attacker can start to search for various strings in the code, such as contact, phone, name, or other information. You can easily correct this vulnerability through minor changes in the structure of Web page design and the thinking that goes into the coding of a Web page, keeping all references to company information out of the page. So even though it's a good programming practice to document your program with comments, in Web design this can be a very bad idea. Keeping the normal programming comments such as "Programmed by John Smith at ABC Company, Telephone 444.4444" out of the Web page can reduce the amount of social engineering that

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:5

might take place. Comments such as "This is where the login authentication takes place" is not something you would want to flag for an attacker. The Web page should use soft links, not hardcoded links with IP addresses or even the full path if it can be avoided. You don't want to give the attacker any freebies with your Web pages to aid his attack on your DMZ.

One of the latest approaches to network security is actually focused on the hosted application itself and has become very important to any company with a server on the Internet. This approach is to run all program code through a quality assurance program to tighten up holes such as buffer overflows and SQL injection attacks. One example of such a program is QAInspect by SPI Dynamics, found at www.spidynamics.com. Within the DMZ, having secure code is of the utmost importance, since there is often interaction between a DMZ host and outside or untrusted users. Virtually all the vendors of the various operating systems and applications have the same issues of buffer overflows. Some vendors have more flaws than others, but in the end, one flaw is one too many. For readers who might not have a good grasp of just what a buffer overflow is, let's take a high-level look at how a buffer overflow attack works.

In programming, various functions are performed in random-access memory (RAM). The program and the operating system allocate certain amounts of RAM for functions such as buffers or temporary data storage areas. This data storage area is generally referred to as the stack. For example, as new data arrives on an Intel stack, existing data is pushed down the stack. Data is retrieved or "popped off the stack" by the use of pointers. A clever attacker can circumvent how the application normally runs by forcing either the application or the operating system daemon or process to error out by exceeding the parameters the application originally set for the buffer space. When the application has an error condition in the buffer space such as exceeding the string length for an input and the subsequent error handling code is poor or nonexistent, a few different results can occur. One is that the application crashes and often-times unceremoniously dumps control over to a command prompt with the permissions assigned to the crashed application. Another classic example of this result is where a long, strange string is sent to a Microsoft IIS Web server and the Web server crashes, running the attacker's code of choice. This particular attack example is based on Microsoft's MS02-065 security bulletin pertaining to buffer overruns in Microsoft Data Access Components (MDAC). The attacker can literally move the stack pointers to run code of their choosing by overflowing the allocated buffer space. Now when the buffer is overflowed and application pointers are popped, they will reference the attacker's code. This new piece of attack code will then be executed. As software complexity increases with each new release, the odds of poor programming practices contributing to these exploits increases.

Correcting this type of security breach will be a multidiscipline effort with the programmers, network engineers, and security officers all having to work together. For example, with distributed applications that use "personal" versions of SQL on the desktop, having bad code across the enterprise will certainly cause problems. All we need to do is to look at the SLAMMER worm attack and understand that the Microsoft SQL servers were not the only things at risk—every desktop that used the Microsoft SQL engine as part of a distributed application was at risk. Now instead of one or two doors for the attacker, he has many doors

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:6

to choose from. So the threat is not only in the DMZ with the SQL servers; it migrates to the desktop. The bar for keeping the DMZ secure has just been raised a significant amount. In this chapter, you will learn how important it is to make certain that publicly accessible DMZ hosts are locked down and secure.

Intrusion Detection 101

Now let's look at our "burglar alarm," or IDS. There are two major types of IDS system. The first is called *host IDS*, or HIDS, and this is where the IDS will run on the host system, or quite simply, on the host itself. The HIDS system uses an agent to monitor the integrity of the local host by watching for changes in various files such as the registry, log files, system files, kernel files, or other specified files. When a change is detected, the agent sends out an alarm, alert, or traps via pager, SNMP, or some other notification method. One well-known HIDS system, Tripwire, is available in an open-source version or a commercial version.

The second major type of IDS is called *network IDS*, or NIDS, and consists of one or more remote sensors with a management node. The NIDS monitors traffic at wire speed and examines each packet in real time to look for a pattern match. One problem that should be readily apparent is trying to have the sensor keep up with today's gigabit speeds on the wire. Some NIDS systems offer different ways to filter out the massive quantity of traffic you would run into by examining every packet.

Some of the significant characteristics of HIDS and NIDS systems are:

- HIDS:
  - Detailed logging
  - Better at detecting unknown or 0-day attacks
  - Fewer false positives
  - Each host needs an agent installed and configured
- NIDS:
  - Ease of deployment (although this can vary with network, hardware/software, and budget constraints)
  - Cost is generally lower (dependant upon architecture and chosen device/sensor)
  - Managed by a single director or management station

Furthermore, HIDS and NIDS can each be broken down into 2 major sub-types. These are signature-based (requiring ongoing updates, much like your anti-virus application) and anomaly-based (requiring an initialization period to learn the idiosyncrasies and behaviors of your network). It is very common to also see devices that offer some crossover or components of both sets of functionality.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:7

When you think of intrusion detection, there are four primary types of data sources to consider.

- Network data packets
- Host details and system details
- Application log files
- Targeted data where the IDS looks for tampered data

Each type of data requires different handling. It's up to you to decide which type of data you need to see or track. The data type will drive, to an extent, the type of IDS system you deploy, or perhaps you'll deploy a mix of sensors. The optimal deployment would be multiple NIDS systems and HIDS installed on key hosts or using a HIDS-like Tripwire to watch key hosts. However in the real world there is something called a budget that brings heartburn to network engineers trying to protect their networks. Personal experience shows that with some careful planning of budget dollars and utilizing resources such as open-source projects, you can protect your network very well with a mix of the two technologies. This way you can budget for the brand-name IDS such as Cisco for the primary IDS sensor and use a "free" product like Snort for the remaining internal sensors. While creating additional complexity in your environment, this also provides informational and system overlap. For example, the Snort user community will oftentimes announce new IDS signatures faster than a commercial provider using a HIDS product like Tripwire to protect the server farm and infrastructure such as the routers and switches you have an effective layered defense that starts from the outside to the DMZ and down to the internal network resources. Given the quality of open-source applications such as Snort, there is no longer any excuse not to have an IDS system or even multiple sensors deployed.

As mentioned in the introduction, there are two primary types of intrusion detection. There is a third method called *network node intrusion detection*, or NNIDS, which is a hybrid of HIDS and NIDS. NNIDS is an agent-based system, but it views packets and runs an analysis against the packets. The difference between NIDS (where all packets are analyzed) and NNIDS is that NNIDS only looks at the packets destined for the host machine it is installed on and protecting. This allows the NNIDS to run very fast. The downside is that you have to install an agent on each host you want protected in this manner. You will see crossover territory with this type of IDS setup that is also referred to as "target-based" intrusion detection.

NOTE

One of the best-known names for looking for tampered data is a product called Tripwire, which can be found at www.tripwire.org. For a poster that has exploits mapped to type and source, go to www.tripwire.com/files/literature/poster/Tripwire_exploit_poster.pdf and find a plotter to print it. (Watch the caps in that URL, the site is picky about the uppercase "T".)

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:8

An IDS deployment has two parts. One part is called the *sensor*; the other is the *management station* or *director*. The sensor is connected to the network and captures a copy of each packet or "sniffs" the packets as they go by. The management station controls the sensors and generates the various reports and alarms from a database of alarm data sent to it by the sensor. The sensor generally has two NICs, one attached to the target network and without an IP address. The lack of an IP address helps protect the sensor from attack. An even better way is to use a network tap on the target network to attach the sensor's monitoring interface. The tap allows you to monitor a full-duplex connection and prevent someone from attacking the sensor itself.

In Figure A.1, we see a basic network tap schematic and the data flow between Network A and Network B. The tap is in between the two networks and provides a copy of both the transmitted data flow and the received data flow. The tap can be passive or active and allow the IDS to send TCP resets to an attacker.

Figure A.1 Common Network Tap

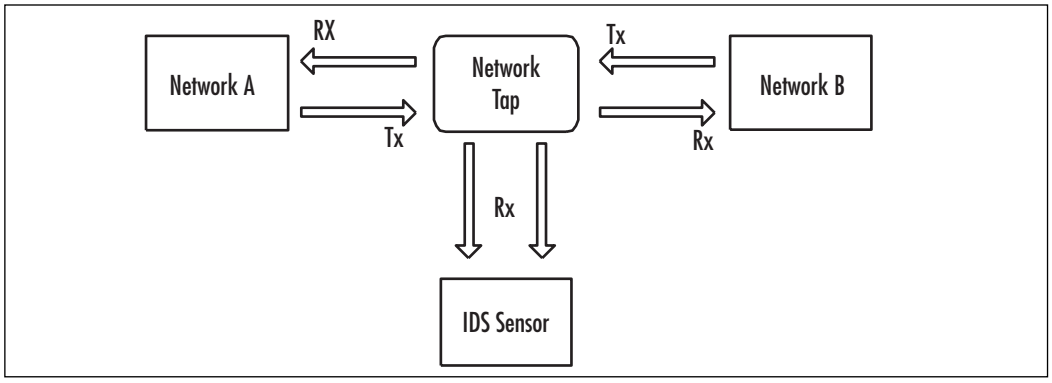

The second NIC is attached to the trusted network and allows communication between the sensor and the management station. The classic deployment of an IDS has the unnumbered port on the outside of the firewall and the management port on the DMZ or on the inside network. An IDS sensor can also be configured with the unnumbered port on a vendor segment and the management port on your trusted network. With enough resources, you can have multiple IDS sensors deployed on the network at the same time watching different areas of interest and each using a different rule set.

An IDS system works by using a sensor to see or "Snort" the raw data off the network or host. Then the IDS engine applies rules to the data to see if the data matches either normal activity or suspect activity, as shown in Figure A.2. Once the match is made, the IDS sends an alert or writes the alert to a log file. The data that the IDS sensor looks for can be quite varied. It can be IP addresses, port numbers or types, strings in the packet payload, or certain combinations of hex with a particular offset in the packet. Due to this flexibility, the IDS sensor can see things that virtually nothing else on the network can see. If an attacker

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:9

launches a port 80 HTTP attack, the firewall will most likely miss it, because it appears to be normal port 80 HTTP traffic. The IDS sensor, on the other hand, sees the possible attack, given that the proper rules have been enabled. The sensor sees the attack because the sensor looks at the entire packet and finds the match of the attack string being used.

Figure A.2 Block Diagram Describing the Way IDS Works

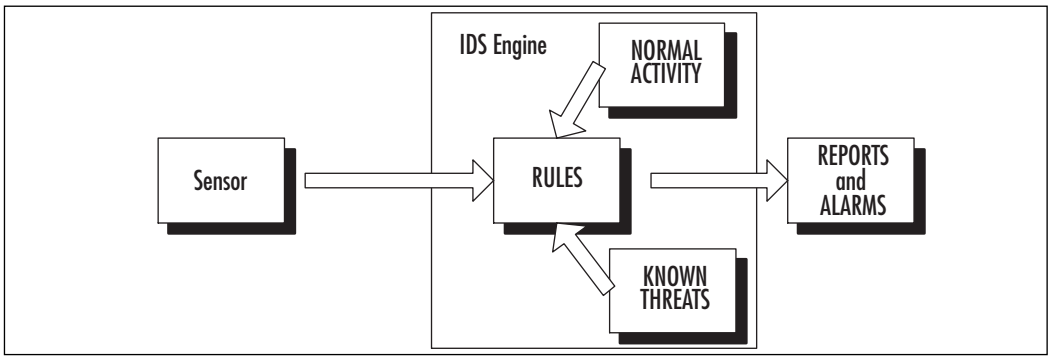

In order to effectively use IDS on your network, you first need to understand how TCP/IP works and the so-called three-way handshake. This three-way handshake is how the classic TCP connection is established between two hosts. The bad news is that an attacker can take advantage of this handshake and manipulate its normal functionality in an attempt to trick routers and/or hosts into giving up information about themselves. For example, different operating systems will respond in various ways upon receiving an initial SYN packet containing data. We discussed how this could be done in some detail in the last chapter. In this discussion, the question is, how do we find this anomalous handshake taking place? We can track this type of activity with an IDS system and the correct signature for the IDS sensor to use. In Figure A.3, we see the normal three-way handshake for TCP/IP.

Figure A.3 TCP/IP Three-Way Handshake

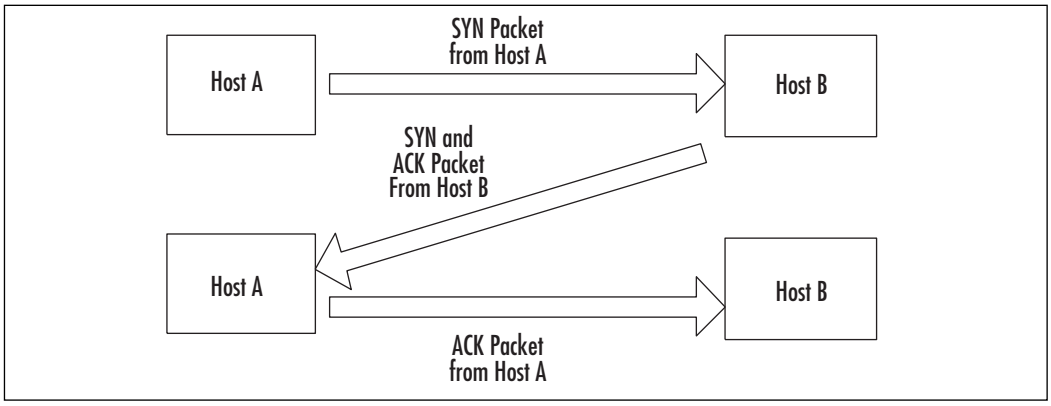

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:10

Designing & Planning…

Where to Put the IDS Sensor?

Many factors drive just where to place an IDS sensor. In the end, cost normally wins out. With that in mind, let's look at the most common places to use one or two IDS sensors. The first is to place the sensor right behind the perimeter firewall and LAN router. This is the network choke point for all traffic that passes in and out of the private network to the outside world. The second choice is to place the IDS sensor within the DMZ. This allows you to watch the exposed servers that face the outside world and look for attacks starting there. Since most attacks are directed against DNS, mail, and database servers and the like, this placement can be of high value. Remember, the IDS is observing traffic, oftentimes looking for attack patterns against a signature database, regardless of whether you are using HIDS or NIDS. The point here is that if you place the IDS correctly, you will be able to observe malicious traffic quite easily, but if you misplace it, you might not capture or analyze anything of importance. The third and somewhat "argued over" placement is between the firewall and the border router. In many cases the border router is the gateway facing your network at the ISP side of the Internet connection.

This style of deployment directly exposes the sensor to the Internet, so care is needed in configuring the sensor so that you don't add a new security problem or issue when securing the DMZ in this manner. This is a perfect example of when to use a network tap, which we discuss in depth later in this chapter.

However, this placement gives a wealth of information as to what attacks are being launched and which actually make it through the firewall.

It is not enough to know the three-way handshake breakdown to effectively use an IDS system. You must also understand, at the very least and at a high level, the various attack methods and how you can make or modify signatures to protect your network against attack. As we saw in Chapter 15, all attacks are not apparent on the surface. Some are written into an HTML header or keyed in at the user ID and password prompt. A good signature will look for these types of attacks. Once the IDS is deployed, don't make the mistake of thinking your job is done—the job has just started.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:11

Configuring & Implementing…

How to TAP for the IDS Sensor

There are a few ways to get the IDS sensor on the wire. In today's switched network designs, it becomes harder to get the IDS sensor to see all the needed traffic due to the way a switch functions when it forwards frames. In simple terms, a switch is really a collection of bridges and operates at Layer 2 of the OSI model. Each bridge or port is its own collision domain, but all the ports share the broadcast domain, since the broadcast domain is set by Layer 3 in the OSI model. The switch looks at the frame and decides which port or "bridge" the frame needs to be sent to based on the destination MAC address. Unlike a hub, where all ports see all frames on the switch, only the port that needs to see the frame will have that frame sent to it unless the frame is a broadcast or the frame has been flooded to all ports. This brings us to our deployment problem with the IDS sensor. Since the IDS sensor needs to see all traffic, what shall we do?

We have a few options, depending on the type of switch and configurations. The first option is to configure the switch for port spanning or mirroring. This would copy each frame and send it to the port the IDS sensor was attached to. One issue to be wary of with this solution—the mirrored port might not be able to handle the aggregate traffic load.

An alternative choice is to use a device called a *network tap,* which is a T-type device that can allow the IDS to sit in the traffic flow from a bidirectional perspective, which can help lower false positives. In some cases the tap can be configured for receive (Rx) only for high security or bidirectional, where the IDS can send RST to clear connections transparently.

Now that you understand the underlying concepts of an IDS let's look at how to deploy an IDS in a DMZ-based network. In this chapter we look at how to set up an IDS in a DMZ utilizing popular vendor-based solutions, including Cisco-based IPS and the freeware program called Snort.

Deployment of an IDS

There are six basic steps for the successful deployment of any IDS system:

1. Determine what network traffic you need to monitor with the IDS to protect assets.
2. Connect the sensors to the network where previously identified traffic will flow.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:12

3. Run the sensors for a week or so with the default configuration. This allows you to build a baseline of alarms and traffic for analysis. This tool helps you decide which alarm is real and which is not and why not. For example, is there a network management software package on the wire that acts like a port scanner? Once you have analyzed the baseline traffic, you can fine-tune the signature-set to reduce the false positives and hopefully avoid any false negatives.
4. Analyze the alarm log and weed out the false positives by editing the signatures, editing the severity level, or disabling the alarm for that signature.
5. Implement response actions based on positive alarms such as sending a pager alert, using the IDS sensors to send a TCP reset packet which breaks down any established connection or connection attempt, or by shunning the source IP address via a null-routing statement or firewall rule.
6. Update sensors with current signatures based on the vendor's release of new signatures or as needed by adding custom signatures.

These six steps reflect a fair amount of work and thinking in order to successfully deploy IDS systems. Along with updating signature files, take the time to make a few minor adjustments to help beat the attacker who is looking to defeat the IDS sensor. Take the interface that is sensing the traffic and remove the IP address or build a cable that is receive (Rx) only by clipping the transmit wires. If you are running Snort, be aware that some distributions will run Snort as root. This is very dangerous to your IDS' security and your network security if the IDS is breached. We delve into this topic in more detail when we discuss Snort later in this chapter.

The IDS must be able to see all the monitored traffic and be able to keep up with the data flow. A switched network can pose special problems, since each port of a switch is considered a network segment. A way around this is to use port mirroring or span commands.

Configuring & Implementing…

How to Configure Cisco Port Mirroring/Span for the IDS Sensor

There are many different vendors of network switches, but here we show you how to configure Cisco switches, since Cisco is arguably the biggest player in networking infrastructure. The goal is to have a VLAN or port mirrored to a single port for the IDS sensor to see the traffic. You do need to read the switch documentation because the procedure will vary slightly from model to model. Furthermore, there are some limitations, such as the number of VLANs that can

Continued

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:13

be monitored or how many mirrors or span ports can be configured. For an excellent tutorial on Cisco *span*, see www.cisco.com/warp/public/473/41.html.

For a CatOS (COS) switch such as the Catalyst 4000, 5000, and 6000 series, our example spans VLAN 1 to module 6 port 2 and is Rx only:

```
Set span <source port>  <source VLAN> <direction>
switch (enable) set span 1  6/2 rx
```

For a Cisco IOS switch such as the 2900 series, our example is port Fa0/1, which will be monitoring traffic sent and received by port Fa0/5:

```
; enter destination port
Switch(config)#int fa0/1


; enter port to be monitored
Switch(config-if)#port monitor fastEthernet 0/5
```

Note that the Catalyst 2900XL and 3500XL do not support *span* in receive direction only or in transmit direction only. You need to take care not to create a spanning tree loop (STP loop). For example, an STP loop is created when a piece of hardware in a layer 2 environment fails and causes the spanning tree protocol in turn to fail. In a physically redundant environment, this creates a loop. There are some additional caveats for the 2950 and 3550 switches from Cisco. Catalyst 2950 switches can have only one *span* session active at a time and can monitor only source ports, not VLANs. Catalyst 3550 switches can support up to two *span* sessions at a time and can monitor source ports as well as VLANs.

A third option is to use a network *tap*, a hardware device that splits a full-duplex connection into two separate streams. The same tap will in many cases only allow the traffic to be sent to the IDS sensor, although some taps, such as SecureNet IDS taps, can allow the sensor to send TCP resets back out into the datastream. The taps prevent a scanner from even seeing the MAC address of the NIC being used as the sensing port on the IDS sensor. This is an ideal method of attaching the IDS sensor to the DMZ, where attackers are certainly going to run various scans against the IP range in the DMZ. It's hard for the attacker to defeat something that they do not know is there watching them. There are several manufacturers of network taps. A couple of the bigger names are Finisar at www.finisar.com and SecureNet at www.intrusion.com. The basic network tap is a small box with three interface jacks on it and a power supply attachment. The tap itself can be for Ethernet, asynchronous transfer mode (ATM), and use copper or fiber optic interfaces. Inside is a circuit that will split the datastreams, and in many, if power is lost to the tap, failover to a pass-through configuration to keep the network connection from being broken. There are special IDS taps that block outbound traffic but allow the IDS to send a TCP reset packet as part of your layered DMZ defense. The taps can be rack mounted or they can be intended just to sit on a shelf in the rack.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:14

At this point you might be thinking, "You just showed me how to monitor my traffic for free, and now you tell about something to buy. Why would I want to spend any money on this?" That's a fair question, but like many things in life, each of these tools has its place, and neither tool will do it all. Port mirroring does have the advantage of being cheap. It also might be able to view all the traffic, depending on the switch, and it certainly is easy to configure and use. However, the downside is that if you try to mirror a heavy traffic load, the switch might not be able to buffer all the mirrored traffic and it might start to shed the excess frames.

A much more subtle issue is that when the packets go through the buffer, the timing is reset so you no longer have an accurate time reference point. On some switches, when the mirror is set up, it resets the ports, which, on a live network, can be problematic if you need to insert the sensor into the network during uptime. As we saw in the sidebar configuration notes, some switches will not allow you to set the data flow on the mirror to be Rx only, which can cause serious network issues if a mistake is made with the configuration. With a network tap, many of the "cons" fall by the wayside and no longer will give you heartburn. In the end, you need to look at your network and the infrastructure and make an informed decision as to the best solution for you and your organization's needs.

In order to see the traffic on the core of the network, it would be best if there were a way to see the traffic on the core switch backplane. Cisco has done just this with an IDS sensor module or blade for the Catalyst 6000 series switch. The Catalyst 6000 IDS module uses an embedded version of RedHat Linux and sees all traffic on the backplane of the switch. The IDS module uses two methods of capturing traffic to analyze. The first way is to use the switch's *span* feature that many of us use to capture traffic for a sniffer. The second is to use a VACL (also called a VLAN ACL). The VACL gets around the limitations of how many span ports you can have on a switch. One of the biggest benefits of the IDS module is that it can capture multi-VLAN traffic, whereas a PC running Snort cannot. The IDS module can capture data at the full 100Mbps and a 64-byte size packet. This translates into the fact that the switch can capture more than 100Mbps of traffic when the packet size exceeds 64 bytes, which your network most certainly will. The downside is that to implement the IDS module requires a very specific Catalyst switch configuration, which does not come cheap. In addition, the IDS module cannot track IP addresses or perform IP blocking the way the appliances can.

Figure A.4 shows the classic deployment of an IDS system when you have a single sensor. The IDS sensor watches all the traffic on the Internet or unsecured side of the DMZ, and then you can compare what is getting through the firewall for possible corrective action.

One of the possible corrective actions is to have the IDS tell the firewall to shun or block an IP address of an attacker. You would see the port scan starting on the Internet side, and then the IDS sensor, either directly or through a management station, would tell the firewall to block the IP address of whoever was scanning the firewall. In the case of a worm or HTTP attack, the IDS sensor would see the match of the signature and could have the firewall block the address, or the IDS could send a TCP reset to break the IP connection between the attacker and the Web server, SQL server, or other host. This type of deployment

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:15

complements the protection that the firewall gives to the hosts in the DMZ. The firewall will pass a port 80 attack since it appears to be normal HTML traffic, so this type of deployment works well and fits into the layered defense model of network security. The downfall to this strategy is that now your attacker is aware of the fact that you are monitoring traffic and performing packet resets. Be wary of an attacker being able to use your defensive measure against your own networks! One favorite spoof attack is to trick the IDS into resetting packets that appear to originate from valid IP addresses.

Figure A.4 Classic IDS Deployment

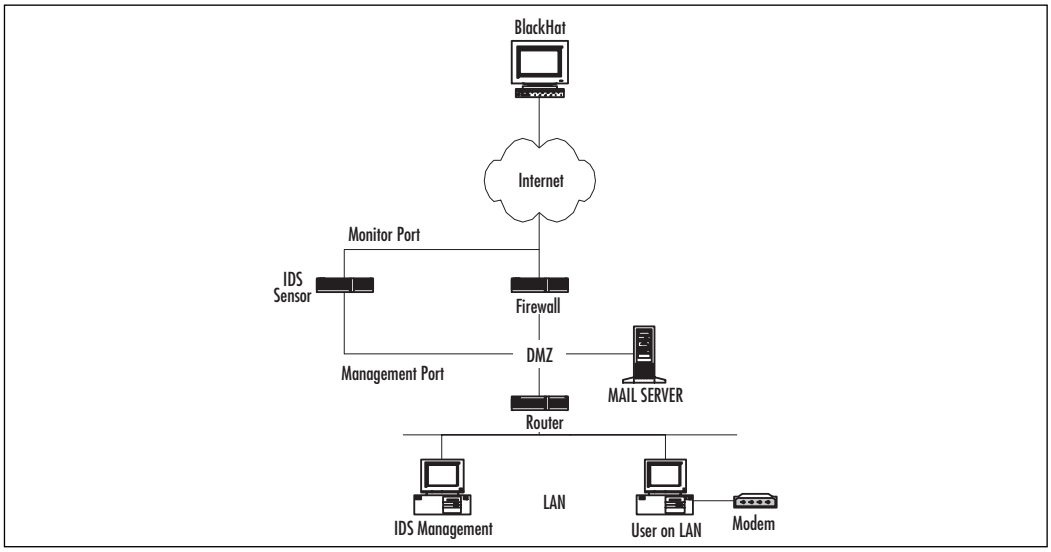

In Figure A.5, we see the results of not having an IDS sensor watching over the DMZ traffic. An attacker uses the Microsoft IIS and PWS Extended Unicode Directory Traversal Vulnerability attack to get past the firewall and, from the Microsoft IIS Web server, get a directory listing of drive C:.

Figure A.5 IIS Unicode Attack Via Port 80 and HTML Results

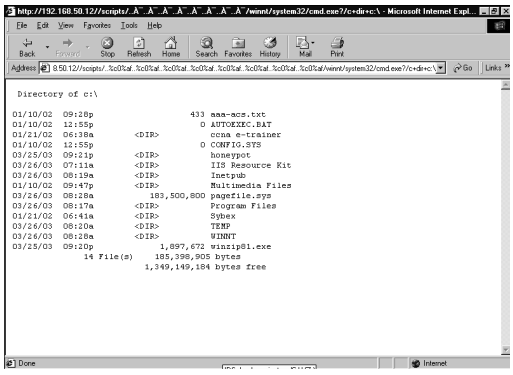

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:16

Many firewalls would not have even seen this attack coming, but many IDS would have spotted it by the well-known Unicode signature. This successful attack has now given the attacker several pieces of information: He knows for sure that the server is an IIS box, he knows the directory structure, and by digging deeper, he could find out if the default IIS files are still there for even more exploits.

In Figure A.6 we see a screen from an Etherpeek capture of a Unicode attack in progress.

Figure A.6 Trace of HMTL Code Sent with Unicode Attack

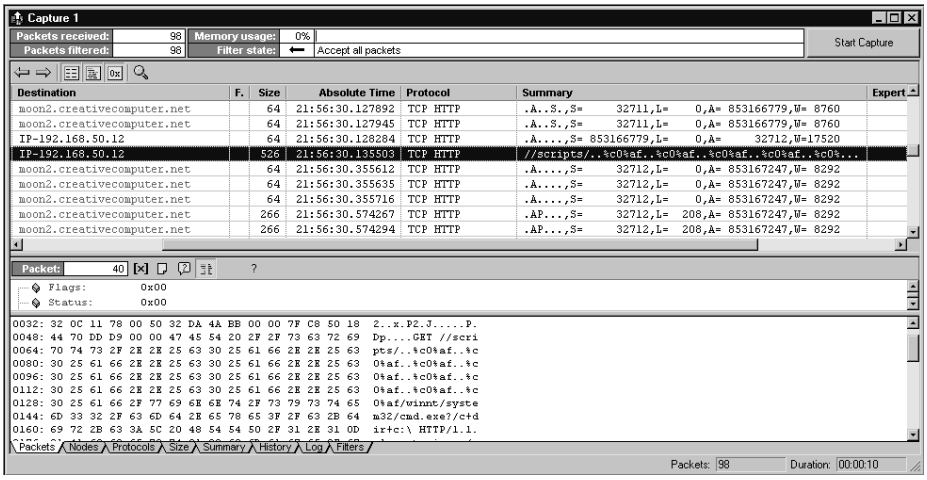

In this capture you can see the signature that the IDS sensor uses to identify the attack, even though it is buried in what appears to be normal HTML traffic, according to the firewall. The IDS scans the payload of the packet and matches the Unicode attack.

By contrast, if the IDS is placed on the inside leg of the firewall, the IDS sees only the traffic permitted by the firewall, and you will need an IDS for each leg of the firewall. So, if you have two DMZs and a LAN connection, you will need three network IDS sensors or sensing interfaces.

This type of deployment can be rather costly in terms of purchasing equipment, time to install the sensors, and time to configure the sensors. The rule of thumb to remember is that if you cannot see the traffic with a sniffer, the IDS will not see it, either. Knowing this, you might want to download Wireshark, a high-quality and free protocol analyzer, set it up, and run a test for a couple of weeks to baseline the traversing traffic. Based on the results of the captures you run, you can see if it is truly crucial to set up an IDS. Another reason you might want to test the traffic is to see if you do in fact need to *span* or mirror the ports on your switch. In this chapter, we look at how to *span* ports on a Cisco switch, where it might be needed. Consult your switch vendor's documentation to determine how to properly configure your switch for mirroring traffic between ports.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:17

Normally there are two interfaces on an IDS sensor. One is for management, and the other is for sensing. For example, on the Cisco 4240 IDS, the built-in NIC is the management interface and the additional NIC is the sensor interface. Snort can be used with a single NIC, but it is preferable to have two NICs installed.

When you're trying to use the *span* or port-mirroring functionality, remember that the commands differ between the Cisco product lines and those of other switch vendors. Some implementations limit the number of span ports you can have, and some limit whether you can monitor multi-VLAN traffic.

When you're tuning the sensor to reduce false positives, it can take a lot of time to sort out whether the alert is real or not. An application such as What's Up Professional from www.ipswitch.com, which uses ICMP *ping* sweeps of the hosts to verify they are responding, can trigger the IDS sensor, since it looks like a *ping* sweep probe. You can edit the signature, filter out the traffic based on source or destination attributes, or go ahead and log it but ignore it.

Typical responses to the sensor having a match of a signature and traffic are:

- IP session logging
- Sending a TCP reset (sometimes referred to as "session sniping")
- Shunning, blocking, or null routing of the offending IP address

Remember that you need to be sure about the positive match before you enable a response. For example, blocking on a port scan could prove embarrassing when the scan comes from your network manager trying to discover nodes on the network. On most IDS systems, the response can be enabled on a per-signature basis. Logging the IP addresses, although it doesn't provide much protection, does provide a way to build trends and to trip the alarm or notification. *Shunning* or blocking can buy time for you to enable a better defense against the attack. Some IDS systems can actually work with the routers and firewalls to dynamically adjust access lists and/or routing tables based on matches to the various signatures.

Repelling the Hacker

One of the most common tools attackers use is Nmap. Nmap can be found at www.insecure.org/nmap/. This knowledge helps you defend the DMZ, since you can look for certain characteristic signatures from Nmap. For example, if the attacker uses the Nmap –sS or "stealth" SYN switch, the pattern can be matched from the IDS perspective, since they are crafted packets. If the attacker uses the older style of connect scanning, it is harder to tell if it is from Nmap, since the IP datagrams are built by the host operating system. You can still find clues based on Nmap itself. If you see an all-port scan and then a second scan of only the open ports, you can be pretty sure someone is running Nmap against your host. Keep this in mind as you check for attacks on your DMZ hosts. To make certain nothing is open for an attacker to see you should run your own scans against your network. Most IDS signature sets will contain appropriate rules for finding these types of packets. One source for

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:18

reviewing the latest Snort signatures (including those related to Nmap) can be found at the Bleeding Edge Snort website at www.bleedingsnort.com.

Another tool to consider along with an IDS system is the honeypot. You can find an excellent source of material about honeypots and the Honeypot Project at www.tracking-hackers.com. For readers who have not heard of honeypots, think about the children's book character Winnie the Pooh and how he is always getting into the honeypot for the yummy honey inside. This is the premise of a honeypot server. It presents itself as a juicy-sweet and irresistible target to the hacker. Once the hacker has "tasted" the honeypot, he keeps coming back for more, and you get to capture it all. For the network administrator or security officer, the real plus is that a honeypot is very cost effective. It doesn't have to be a megaserver to be effective; a leftover P150 with 128MB of RAM can work very well.

The honeypot server can sit in the DMZ and perform an interesting function. It does nothing at all; it's the laziest server possible. In reality, the honeypot provides a tempting target for aspiring hackers to start probing, and since the honeypot does nothing, when there is a probe against port 80 or port 25, it can be assumed that the probe is real. This greatly reduces the noise and false alarms that can occur in a busy DMZ on the traditional IDS. The honeypot can be anything from a simple application to a very sophisticated design, depending on the business's needs and wants. It can provide just an early warning that probes are taking place, or it can act as a logger of all keystrokes and other information. You can use this information to build custom filters or signature files for the IDS. Let's look at how to create a basic honeypot to be placed on your DMZ segment and examine the types of attacks and scans that can be used against it, as well as what to do when you catch someone probing your network services and resources. We start with Figure A.7, which shows the basic deployment of a typical honeypot on the DMZ.

Figure A.7 Basic DMZ Deployment of a Honeypot

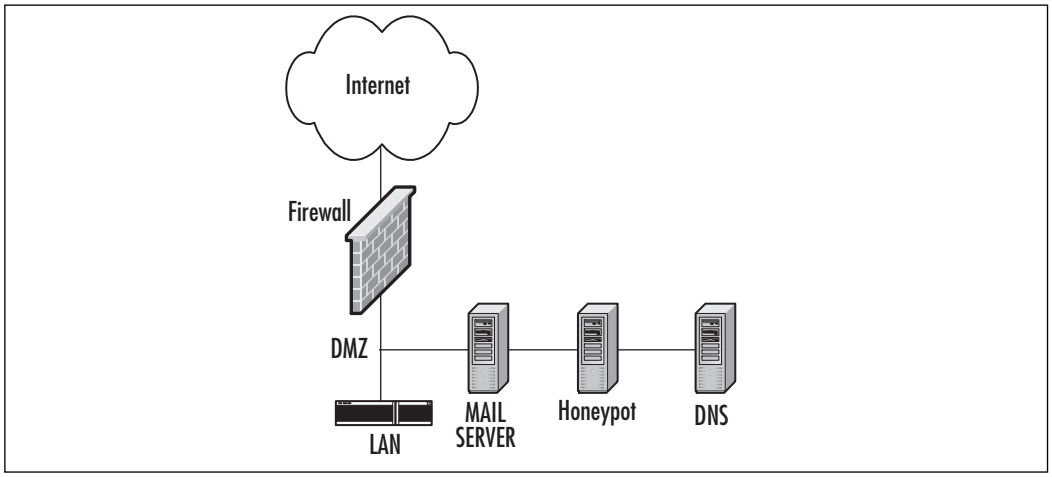

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:19

In this case, as we described earlier, the honeypot is not part of the network, so there should be no connections to it from legitimate network traffic. If there is a network connection, it can be assumed that a hacker has found it via some type of reconnaissance port scan or possibly from another form of passive attack, such as eavesdropping, where the name of the game is information gathering. The IDS might or might not have detected the port scan, so the honeypot has provided a backup method to alert you to the potential of a hacker knocking at the door. To make this concept clearer, a honeypot is not used *instead* of an IDS but *in addition* to the IDS sensor. If something is missed, you could record some poking and prodding on your honeypot, thus clueing you into the possibility of someone "hunting and pecking" your Internet-facing hosts for open services and known vulnerabilities.

One of the best honeypots to start with is Back Officer Friendly (BOF), which is a parody on the remote-control Trojan called Back Orifice. You can find BOF at www.nfr.com/resource/backOfficer.php. It emulates several ports such as Telnet (23), SMTP (25), FTP (20 and 21), and others. When probed on these ports, BOF sends back fake replies and logs the interaction, allowing you to see if someone is knocking on your door. When it comes to your DMZ hosts, you must assess the need for this type of security product, because you most likely will be probed eventually.

A more sophisticated commercial honeypot that can emulate 14 different operating systems and has a wealth of logging and tracking functionality is called Specter; it can be found at www.specter.com. Specter is a very sophisticated honeypot program that can appear to be a Windows 98, NT, 2000, 2003 or XP server. It can also seem to be a Linux, Solaris, Irix, or other UNIX style of box. For the kicker, it can appear to be a Macintosh running OS-X. In addition to the 14 different operating systems that it can emulate, Specter will monitor up to 14 ports. Specter has added an SSH trap and can use WHOIS to backtrack the attacker based on the IP address. You can set up Specter with false users, false Web pages, bogus e-mail, and other tantalizing tidbits to entice the would-be hacker to spend far too much time playing in the honeypot.

On the con side, Specter does not emulate IP stacks, so an Nmap scan shows the truth of the honeypot if the attacker takes time out from his excitement to find an open system to even scan the honeypot. Specter can also send its log files to a syslog server or to the local event log. You can also configure alerts that can be e-mailed or sent to a pager. Keep in mind that although Specter is a great honeypot, you still need to secure the OS you install it on. Specter does not secure the OS; that's up to you.

Honeypots in the DMZ

In this section we look at how to set up a honeypot with Specter. Again, it is crucial to your security posture that you harden and protect the DMZ segment as much as humanly possible, and setting up a honeypot offers you many strategic security options for doing so. Let's take a look at how a honeypot can be used.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:20

Configuring a Honeypot for Your DMZ

One of the best locations for your honeypot is in the DMZ of your network. The honeypot is protected to a degree but gives enough back to the hacker to keep his interest away from other hosts on the DMZ. As we said before, what better way to find out there is a hacker than by having the hacker "ring the doorbell." Specter is easily installed, just like any other normal Windows application. Getting from start to working honeypot can happen in the range of 10 minutes. In Figure A.8 we see the basic "Specter Control" GUI when we start to configure it.

Figure A.8 The Initial Specter Screen

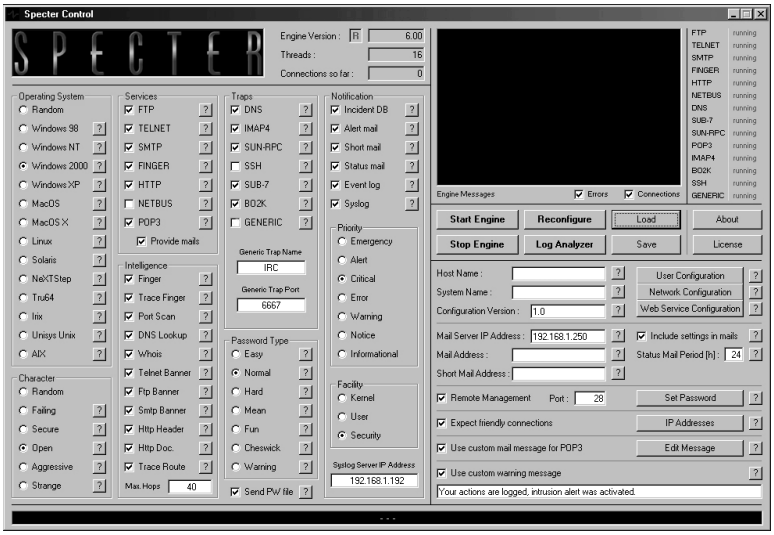

The authors of Specter assume that you know the basics of the operating system you are trying to simulate. Virtually everything is configurable. For example, you can configure how difficult the password should be. There are 7 different setting choices, including easy, hard, or the best one, fun. We suppose that fun depends on one's point of view in this case. You can configure the various banners for things such as FTP or Telnet. You can turn on or off various services such as FTP, Telnet, SMTP, and things like SUB-7 or BO2K. This honeypot is not an interactive honeypot — in other words, the attacker does not get to play with the operating system. This makes a honeypot such as Specter relatively easy to deploy and run as you control exactly what the attacker can do. The downside is that you will not be able to recover interesting things such as root kits, e-mails, toolkits, or IRC chat logs.

To deploy Specter, you need a reasonably powered machine. The published guidelines call for a minimum of 1.5GHz processor with 512MB RAM and Windows XPSP2. Specter uses the normal Windows installer; once the product is installed, just start the honeypot. Each configurable parameter has a question mark (?) next to it; click it to get help if you

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:21

need some explanation of the choice. You can remotely manage Specter, but that service is turned off by default, so you need to remember to enable it. The interface for remote management is virtually the same as if you were on the honeypot as a local user.

Specter offers an excellent remote alert function. This is one of the most valuable benefits of Specter, since the sooner you know the doorbell is ringing, the sooner you can stop the attacker. You can choose to have e-mails sent to you, pager alerts, log information to a syslog server, or have the information sent to the local event log. In Figure A.9 we see a log entry that shows an FTP session starting. The attacker thinks that the password file was transferred to his machine.

Figure A.9 A Specter Log File Screen

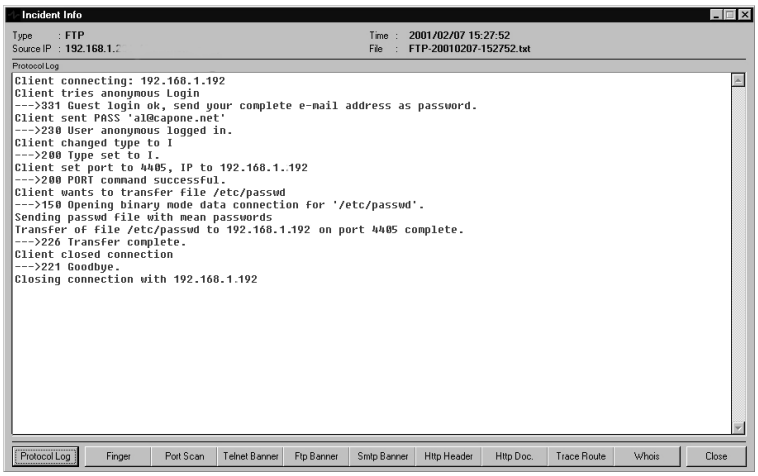

Specter or any other honeypot tends to work best on the DMZ or internal network. We already know that "bad people" are on the Internet, but we want to know who is poking around in our DMZ. For more details on honeypots and their use, see www.tracking-hackers.com for a good selection of resources and articles.

Now that we have learned how to implement a solid honeypot (and how to use it to find attackers), let's take a look at other types of hosts in the DMZ that you need to be concerned about, especially when dealing with HIDS.

Host-Based Intrusion Detection Systems

As we saw earlier in the chapter, HIDS are just what their name implies. The intrusion detection takes place locally on the host that's being protected. This can offer significant value to the intrusion detection process, but it also carries some issues.

Some of the value of HIDS is that the log files can be very detailed. You can see the actual results of an attack; the HIDS does not rely on signature files as much as the NIDS does, which can let the HIDS catch an unknown attack that the NIDS would have missed.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:22

Furthermore, the HIDS can catch an encrypted attack, since it would be decrypted at the host. All this good stuff comes with a price in the form of the possibility of high overhead on the host, the need to secure logging because many attacks will also change the log files, suspect logs if the attack was a success, the need to monitor each host, and possible cessation of monitoring due to a system crash.

The HIDS can watch for many different changes on the host, ranging from files being changed, time stamps, privileges changing, and deleting of data to system modifications. With any HIDS, you need to make sure that the log files have enough space to be stored. This is very important in Windows-based hosts because the default log file size is only 512Kb—far too small for any real security application. Given that disk space per megabyte is so cheap, there is no excuse for not having large log files. The rule of thumb is to have a week's worth of data in the log files. This means that you need to run a few baselines before turning the HIDS on as a production system. You need to determine the size the log files will be with normal traffic then add to the number to account for any unexpected increase in activity such as might happen in an attack.

While we're speaking of log files, we want to discuss configuring the hosts monitored by HIDS with Network Time Protocol (NTP) clocking. If you need to analyze an attack spread over several hosts in the DMZ, you need to have all the log files on the same time standard so you can build an accurate timeline of the event. For legal purposes, the log file's timestamp must be in the Universal Coordinated Time format. Without an accurate timeline, deciphering and prosecuting an attack is much more arduous.

Some HIDS can gather information from hosts even if the software is not resident on the device. This feature can prove very useful in guarding your network infrastructure, such as the routers and switches—for example, if the router log files are configured to track exceptions to the access lists in place. This can provide an alarm that someone is knocking at the door or give insight into how the attack might have started. Even without a HIDS, the syslog functionality alone can be very helpful. With a combination of log severity level and facility levels, you can configure multiple entries on a syslog server such as Linux. For one thing, you could have two PIX firewalls and set up a custom directory for each PIX using the facility option of syslog. For analyzing the log file data, some excellent applications are available. One of the best is called Sawmill (www.sawmill.net). This application can read virtually any log file output and generate some excellent reports from the data. You can analyze your Web server log files, firewall log files, Cisco IDS log files, or Snort log files, just to name a few.

The world of HIDS holds a few major players and many minor players. Before you deploy a HIDS, you owe it to your sanity to do some research on which HIDS is the best one for your network design, risk factors, and budget. A few of the bigger names in HIDS are Tripwire, found at www.tripwire.com; Cisco Security Agent, at www.cisco.com/go/csa; and McAfee Host IPS, at www.mcafee.com. Let's look at one of these products in a bit more detail.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:23

Tripwire

Tripwire started life as a UNIX-based academic source project in 1992. It has morphed into a commercial project for both UNIX and Windows. A Linux version has been released as an open-source project and can be found at http://sourceforge.net/projects/tripwire.

Tripwire, classified as a HIDS, provides an effective method to manage data and network integrity. It does not look at packets, the way a NIDS does; it watches the files and configurations of network nodes. To manage the nodes, Tripwire communicates with the various hosts and nodes on the network using Telnet/TFTP, SSH/TFTP, or SSH/SCP. The information gathered is maintained in a database and used to make a baseline signature for each node. In Figure A.10 we see the basic Tripwire screen showing nodes that are monitored.

Figure A.10 Tripwire Node Manager

Tripwire is managed by an HTML interface, so the administrator can manage it from anywhere on the network. Tripwire's logging function is very sophisticated, with features ranging from basic reporting of alarms to the capability to quarantine evidence for forensic analysis. A very useful feature is that Tripwire can be configured to restore a changed node configuration either manually or automatically. This has some interesting possibilities for managing nodes in the DMZ. If someone were to attack and modify the configuration of a host/node in the DMZ, Tripwire can put back the old configuration as soon as it spots the change. In Figure A.11 we see that Tripwire detected a change between the "startup configuration" of a router and the current "running configuration" of the same router.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:24

Figure A.11 Tripwire Detecting a Change in Configuration

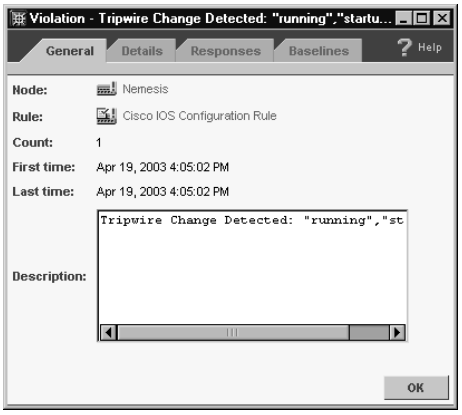

We can configure a rule that has Tripwire put the "baseline" image back onto the router automatically when a change such as this is detected. Of course, you would only want to do this to nodes that do not see very many changes in their configuration, but you can set the rule to be manual in order to prevent Tripwire being overly helpful by overwriting changes that you just made to the node.

HIDS such as Tripwire can watch more than infrastructure nodes. They can monitor various file characteristics such as create time, last access time, user ID, file permissions, alternative datastreams, various Registry keys, and much more. Most HIDS such as Tripwire monitor both Windows and UNIX machines of various flavors such as AIX, Linux, Solaris, and HP-UX.

Ideally you would team the IDS sensors with a HIDS (like Tripwire) for solid intrusion detection coverage and then use a honeypot as the early-warning system — the "doorbell" that we have talked about. Now we are getting closer to the "onion"-layered network defensive posture we discussed earlier.

Saving the DNS Server

The poor DNS server is one of the mostly likely targets in the DMZ for attackers to nail first. The reason is that most enterprise networks have a DNS server located on the DMZ segment, and with the number of attacks that can be launched on it, you will more than likely log some attempts on its life. Given how DNS works, you might think that all is lost and that protecting the DNS server in your DMZ is just about impossible—but nothing is further from the truth.

We can do several things to harden the DNS server against the malcontents of the Internet:

- Use the newest version of BIND, which is currently 9.3.2. You might also want to consider running a different DNS service such as *djbdns,* found at

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:25

http://cr.yp.to/djbdns.html. This service is designed to be a hardened DNS service. Windows 2003 does offer a DNS service, but hardening a Windows 2003 DNS server takes quite a bit of work.

- Keep your patches up to date. Doing so is critical with all vendor products you are using in your DMZ, especially when the hosts deploying these services (such as DNS) are Internet facing and open to scanning by attackers.
- Restrict or authenticate zone transfers to prevent unauthorized release of your zone information. Give no freebies to the wannabe hacker of your DMZ.
- Restrict DNS dynamic updates by using access lists to control exactly who can and cannot get the updates.
- Use the Split DNS architecture – where publicly offered IPs are maintained on a separate DNS server.
- Use the most secure method of updating your DNS server records. This includes limiting who has access to TCP port 53 on your DNS server as this port is used to perform the zone transfers that update your DNS records.
- Make sure the firewall is configured correctly for DNS traffic on both TCP/UDP 53.

The two items that bring the most value to hardening of your DNS servers is the updated version of BIND and the split DNS architecture. The idea of split DNS architecture is pretty straightforward; SANS has a good article explaining split DNS at http://www.sans.org/reading-room/whitepapers/firewalls/791.php. You will have two (or more) DNS file servers. One will be in the DMZ and offer public IPs; the second will be internal on the trusted network. One server will forward requests to the other.

When configuring the firewall, think about how DNS works, and make up access lists to help provide defense in depth. For example, if you have an internal DNS server, the only DNS server that should be asking for any information at all should be the DNS server in the DMZ. So, the ACL should deny all DNS traffic except for that server. Make sure that the attacker cannot use the *dig* or *nslookup* commands to pull up the entire zone. You can accomplish this by using a feature called *xfrnets*, found in BIND versions 4.9.3 through 4.9.5. This feature allows you to edit the /etc/named.boot record and apply an access list to zone transfers by IP address. In BIND 8.1 or higher, this functionality is called *allow-query* and also applies an access list on DNS queries.

We also see that when you run the *dig* utility and ask for the version of BIND, as shown in Figure A.12, you will get the DNS server to respond with the version of bind it is currently running. This gives the attacker a rather large clue as to how to begin to attack the DNS server. We want to hide that information from the attacker. It can be done, but it will take some effort on your part.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:26

Figure A.12 A DiG Query of BIND Version

```
[msweeney@venus msweeney]$ dig @66.75.160.41 version.bind txt chaos

; <<>> DiG 9.2.1 <<>> @66.75.160.41 version.bind txt chaos
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 19939
;; flags: qr aa rd; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;version.bind.                  CH      TXT

;; ANSWER SECTION:
version.bind.           0       CH      TXT     "9.2.1"

;; Query time: 27 msec
;; SERVER: 66.75.160.41#53(66.75.160.41)
;; WHEN: Sat Apr  5 07:08:32 2003
;; MSG SIZE  rcvd: 48
```

As you can see in Figure16.12, BIND happily gave out its version as 9.2.1 when we performed a *dig* query asking the server for the version. This information tells the attacker the version of BIND he needs to find an exploit for.

How do we get BIND to stop being so helpful to intruders? In Figure A.13 we see the simplest way to stop BIND from this kind of helpfulness. We set an options sub-statement in BIND that gives back a string that we define in the version argument.

Figure A.13 A Way to Stop Bind from Handing Out Its Version String

```
options {
    directory "/var/named";
    version "They stole it from us";
};
```

As you can see, with just a few simple fixes, you can deter any potential hacker who might want to compromise your DNS servers. Remember what the Greek philosopher Demosthenes said: "Small opportunities are often the beginning of great enterprises." We keep this statement in mind to resolve our DMZ concerns, where "small opportunities" to lock down vulnerabilities like this could be the beginning of a great (and secure) enterprise DMZ.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:27

Keeping the Web Server Serving

To keep your Web server from becoming a hacker's playground or to prevent hackers from showing off their newest graffiti or Web page manipulation skills, you need to pay attention to some simple rules while addressing security on a Web server in your DMZ.

As with the DNS server, keeping the Web server up to date on patch levels is of the utmost importance. For example, on Microsoft IIS servers, applying best practices hardening guidelines in conjunction with keeping current with Microsoft's patches and hotfixes can prevent virtually all known exploits to date. Although sometimes it feels like there's an endless stream of hotfixes and patches, the fact remains that simply keeping the software up to date is one of the best ways to help keep your Web server from being hacked.

A second simple way to help harden the Web server is to disable *all* unneeded services and don't run too many things off a single box. Don't try to make a Web server also perform as a DNS server to save time and money. You will just buy yourself trouble in the long run. Other simple ways to help keep the hacker at bay: Either disable or harden access to programs like FTP, tftp.exe, command.com, telnet.exe and cmd.exe so that script kiddies end up getting discouraged and leaving your server alone, moving on to easier prey. A hacker will do sweeps of public IPs on the Internet looking for blocks that are assigned to specific companies. If a hacker scans, say, 1.1.1.1 through 1.1.10.255, he could find hundreds of IP addresses that are actively assigned to many systems either accessible from the Internet or sitting within a DMZ segment. If he cannot hack your server, he might very well move on to the next host standing by, maybe with more services or vulnerabilities available.

For the various flavors of UNIX/Linux, you should strongly consider downloading and running SELinux, which can be found at www.nsa.gov/selinux/, or Bastille-Linux, which can be found at www.bastille-linux.org. Bastille-Linux attempts to harden your Linux installation and supports Red Hat, Debian, Mandrake, SuSE, and TurboLinux Linux distributions, along with HP-UX and Mac OS X.

For the Windows IIS server, a quick Google search on "IIS 6.0 Security Best Practices" will land you at Microsoft's best practices guidelines page. An entire chapter is dedicated to Security in IIS 6.0 under their IIS 6.0 Operations Guide.

Additionally, it's an easy administrative tweak to rename the administrator account and disable the guest account on the Windows IIS servers. Although this will not prevent the expert hacker from finding the administrative account, it will deter those less talented hackers, and in the end, that's what we want. We want the hacker to go away to play on and exploit someone else's less protected server.

Never forget about passwords on any Web server, or any server for that matter. Always use strong passwords; choose non-word passwords so that basic dictionary attacks will fail. Use policies to lock out accounts after too many login failures to help discourage the would-be attacker. Log failed attempts and…review your logs!

Now that we have reviewed the reasons that Web services are so vulnerable, let's look at how to put the icing on your DMZ-cake and utilize HIDS on your Web services in the DMZ—the most likely target, and most common target, within the DMZ.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:28

NOTE

Before we move on to applying intrusion detection and logging to your Web services, note that many organizations opt not to have a Web server on site. Instead, they place Web servers in an offsite location, remotely managed, outsourced, or colocated. If you have these types of solutions on your own corporate network, you should still ask your provider if services such as the ones listed in this chapter are applied and used. It is important to know this because often, depending on the provider, these services are not used. Many hacked Web sites are generally located on a third-party Web server on which a company might rent space. If you have your own server or one located elsewhere, ensure that your solution (wherever it is) is secured in the same ways listed here.

Configuring & Implementing…

Securing the E-Mail Relay with Intrusion Detection

In previous chapters, we have seen how a potential hacker could compromise your e-mail server in the DMZ to send spam or use the mail server as a relay. IDS can be a powerful tool in stopping this abuse of your e-mail server. In the following example, we have a scenario of an attacker wanting to DOS your e-mail server in the DMZ. By running an Nmap scan against your DMZ, the attacker has found that you have a Domino e-mail server running. You thought your network was safer than the average network since the e-mail server is not Exchange, but our wannabe hacker is not deterred. He does some research and finds an exploit for Domino that can either crash the server or possibly give him the ability to execute some code. All he needs to do is run a very long *rcpt to* command to start the buffer overflow (bugtraq: 20010123). What our hacker doesn't know is that you are running an IDS via a network tap that hides the IDS from his scanning. Furthermore, you have a rule that looks for this type of attack. The rule, in a nutshell, looks for a string of any case over 800 characters following the *rcpt to* string. So our hacker tries to run his attack against your e-mail server, but the IDS sees the match and then sends a *TCP RST* to the hacker's IP address. This resets the hacker's TCP connection to your e-mail server, and poof! He is sitting there wondering why his connection to your e-mail server just dropped. This is just a sample of how an IDS can be configured to protect a certain host—in this case, the SMTP server—in the DMZ.

Continued

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:29

The same IDS can also watch for an error condition that will alert you to someone who might be trying to relay mail through your server. Again, using Snort as the IDS, we would have a rule that looks for the string *smtp-relay-denied*. A match on this rule could be configured to send a page to alert you that someone is poking around on the e-mail server, or attempting to.

Cisco Enterprise IPS Management

Successful attacks against enterprise networks typically require a substantial effort on the part of the attacker. Many large networks that discover they have been compromised only do so after discovering a discrepancy in activity or log files traversing their network. Once the compromise is known the network staff may backtrack and identify all of the activity that occurred *prior* to the compromise…or they may not. Attacks typically are characterized by three phases of activity:

- Reconnaissance
- Probing
- Exploitation

Reconnaissance involves identifying network address ranges, telephone numbers, performing DNS lookups (both forward and reverse) as well as *whois* searches to identify potential names and accounts to try on various target systems. Probing involves ping sweeps to identify potential targets as well as port scans to identify services active on the target systems. Finally, exploitation of a vulnerability (whether it be a buffer overflow in a running service or simply access through poor password selection) is the culmination of an attack to gain access to the target network.

The probing and exploitation phases require the use of active tools to identify available services and potential exploit targets. It is this activity that Intrusion Prevention Systems (IPS) are designed to identify. By monitoring traffic on the network and inspecting and analyzing packets the IPS is able to determine if a network is under a possible attack. If an attack is identified by the IPS it can issue alerts to network and security operations personnel so that they can respond appropriately to protect vital corporate assets. Additionally, many modern IPS systems can execute response measures on their own accord thus terminating the attacker's connection.

There are significant differences between managing a small handful of IPS sensors (on the order of one, two, or three sensors) and managing an enterprise-wide deployment of sensors. Tuning a single sensor to the traffic on a particular LAN may require one or more days simply for the actual tuning of IPS signatures. Once that has been completed the sensor must be monitored for false positives and for any additional signature tuning required. This can take on the order of a week or more for a single sensor. When new signature packs are released containing additional attack signatures they must be deployed and tuned as

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:30

well. Clearly once the number of sensors goes beyond a small handful then the administrative effort of configuring, monitoring, and updating sensors becomes a significant burden. By using a tool that provides for managing all sensors through a single interface the burden is dramatically reduced. This is where CiscoWorks VMSCiscoWorks VMS and, in particular, the Management Center for IPS Sensors (MC) are meant to provide the greatest benefit. Scalable management of IPS sensors is needed to meet the needs of an enterprise network. The Cisco Intrusion Detection System Management Center is designed to provide the centralized sensor management required to protect large enterprise networks.

Understanding the Cisco Management Center for IPS Sensors

The Cisco IPS Management Center serves four primary functions:

- Logs audit records pertaining to the Intrusion Detection System (IPS)
- Notifies IPS personnel when internal event thresholds are reached
- Manages and distributes configurations to the sensors
- Manages and distributes signatures to the sensors.

IPS MC & Security Monitor

Closely related to the Cisco MC for IPS Sensors is the Cisco Monitoring Center for Security, also known as the Security Monitor. Although the Security Monitor is a separate and optional product, it is often packaged with the MC for IPS Sensors. While the Security Monitor's primary purpose is to receive alarms from the Sensors, the MC for IPS Sensors's primary purpose is to administer and manage the Sensors.

The Security Monitor provides the following functions:

- Event Collection
- Event Rules and Notifications
- Real Time Event Viewing
- Reporting (scheduled & on-demand).

The Event Viewer of the Security Monitor is used for the real time display of alarms generated by the IPS Sensors. While the Security Monitor may be installed on the same host platform as the MC for IPS Sensors, often it is installed on a separate host platform for increased performance.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:31

The MC for IPS Sensors & Sensors

The Cisco IPS Management Center can manage up to approximately 300 sensors. In the example deployment shown in Figure A.14 below the sensor is deployed on the network perimeter or demilitarized zone (DMZ). Inside the protected network is a management host with the MC for IPS Sensors installed.

Figure A.14 MC for IPS Sensors and Sensor

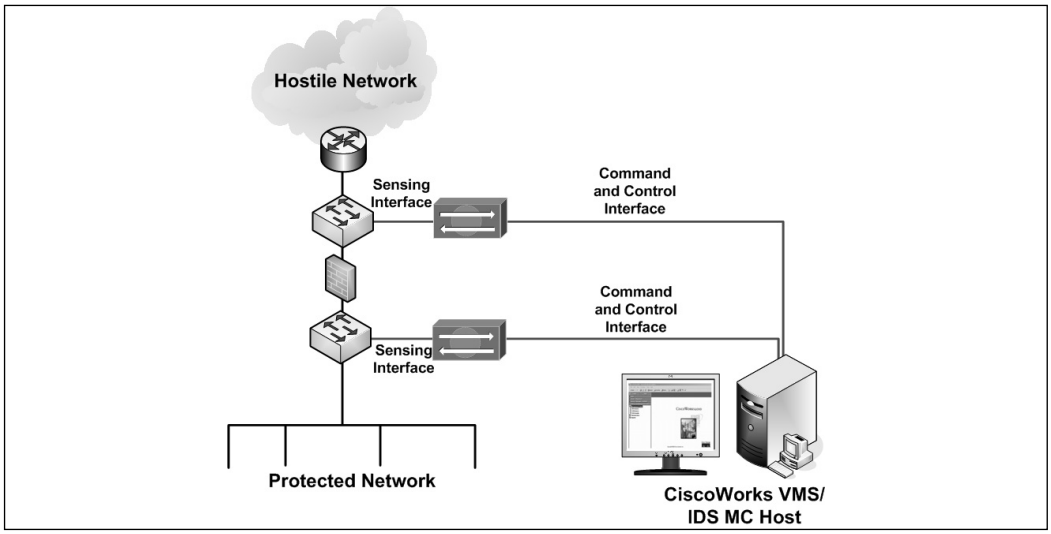

The sensor monitors traffic inside the DMZ between the inner and outer firewall routers. The sensor has two interfaces: a control interface that is connected to the internal network and a monitoring interface connected to the DMZ network. The control interface provides for management and configuration of the sensor. The monitoring interface, operating in promiscuous mode, passively listens on the DMZ segment. When the sensor detects suspicious network traffic on its monitoring interface, it will send an alarm or event to the Security Monitor via the control interface. Through this same control interface, the IPS Management Center manages the sensor and updates its software versions and signature releases. The sensor uses the control interface to enable blocks or shuns in routers or PIX firewalls. When the sensor uses a TCP RST as a countermeasure against an attack it sends the TCP RST packets out through the monitoring interface.

MC for IPS Sensors & Signatures

IPS sensor signatures are the representations of patterns that have certain characteristics of various attacks and other activities attackers may use against a network. The patterns or signatures will be used by the Cisco IPS sensors to detect malicious traffic and to act on it.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:32

Upon detection of a suspected attack or reconnaissance the IPS sensor can send an alarm to the Security Monitor or attempt to intervene through the use of shunning, blocking or TCP resets (RST). The MC for IPS Sensors provides many administrative services with regards to the maintenance of signatures. The MC can be used to enable or disable various signatures based on the administrator's determination of whether they are relevant to the network being monitored by a given sensor. Additionally, the MC for IPS Sensors provides for the capability to define custom signatures that may not be part of the normal signature pack distributed in CIDS software or signature updates. This capability allows security staff to add to the sensor signature database. Managing, updating and distributing these signatures are key administrative functions of the IPS Management Center.

MC for IPS Sensors & Security Policy

From an enterprise perspective, it is important to note that sensor and signature management are merely tools used to implement your Corporate Security Policy. This policy will determine how you deploy your sensors and what signatures you will need.

Designing & Planning…

Cisco Security Wheel

Network security methodology has become increasingly important in driving the overall security of a network. The "old-world" enterprise network security philosophy calls for the development of a security policy first and only then are security products deployed. Once the network is secured, it may be inspected once in a while to identify any potential issues and security incidents are handled on a case-by-case basis.

The Security Wheel is a concept whereby the corporate security policy forms the hub around which all network security practices are based. This methodology evolved as an alternative to the traditional approach to network security. As in the "old-world" approach, the security policy is developed first and then the network is secured according to the policy documents. Also, as in the "old-world" approach, the network is monitored for events and any incidents are handled appropriately. However, where the Security Wheel goes further than the traditional "old-world" approach is that the Security Wheel calls for the testing of network security and the results of that testing then feedback into the security process to manage and improve the state of the network's security posture. This concept is shown in Figure A.15 below.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:33

Figure A.15 Cisco Security Wheel

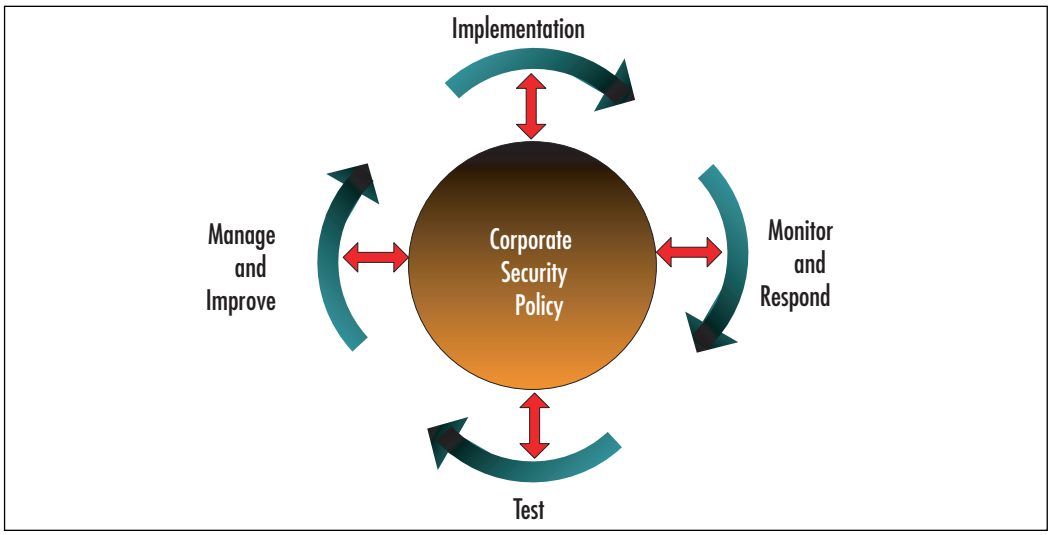

The security policy must clearly state the organization's stance and objectives with regards to security issues. Typically a security policy is not a single document but a group of documents that provide a high-level overview of security implementation in the network. The policy should document resources to protect and identify the network infrastructure and architecture in general. Finally, the security policy should clearly identify any critical resources that require additional protection. Intrusion detection can be seen as an *extension* of the network security policy. In many respects, IPS can be considered the *enforcement* of that policy because it provides a continual audit of the network traffic. An in-depth discussion of the development of a security policy is beyond the scope of this chapter as well as this book. For a more detailed discussion of security policies and how to develop them, please refer to the bibliography at the end of the chapter.

Installing the Cisco IPS Management Center

The Cisco MC for IPS Sensors is a component of the VPN/Security Management Solution (VMS) which, in turn, is part of the CiscoWorks VMS software package. The VMS software suite includes additional components such as CiscoWorks VMS Common Services which provides for user roles and MC access privileges to be defined as well as data storage. Additionally, it provides data storage for CiscoWorks client applications that use its services. Other components of VMS include the PIX Management Center, VPN Router Management Center, VPN Monitor, the Cisco Security Agent Management Center, and the Security Monitor. This discussion will focus primarily on the installation of the MC for IPS Sensors however it is often helpful to understand the combined installation requirements of the required VMS bundle.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:34

Server Hardware Requirements

CiscoWorks VMS can be installed and operated on either a Windows 2000 Server platform or a Sun Solaris platform. The hardware requirements for CiscoWorks VMS and VMS are specified below in Table A.1.

Table A.1 Server Hardware Requirements

Component	Minimum Requirement
Hardware Pentium CPU	IBM PC Compatible with 1GHz or faster Sun UltraSPARC 60 MP with 440 MHz or faster processor or Sun UltraSPARC III system (i.e. Sun Blade 2000 or Sun Fire 280R)
Operating System Advanced Server (Service Pack 4)	Windows 2000 Professional, Server, and Solaris 8 with the following patches: 108528-13 108827-15 108528-13 108827-15
Additional Software later	Microsoft ODBC Driver Manager 3.510 or N/A
File System	NTFS UFS
Memory	1GB (minimum) experience shows 1.5Gb to 2Gb is more practical
Virtual Memory	1GB (minimum)
Hard Drive Space	9GB (minimum) This will increase depending what reporting and logging is enabled. Be sure to monitor the log file size.
Java Version	Sun Java plug-in 1.3.1-b24

The VMS product suite should not be installed on a Microsoft Windows Server system that is a domain controller or a terminal server. The remainder of this chapter will focus on the installation of the MC for IPS Sensors on a Windows 2000 system. For additional information regarding the installation of the MC for IPS Sensors on a Solaris 8 server, please refer to Cisco's website http://www.cisco.com.

CiscoWorks Architecture Overview

The MC for IPS Sensors architecture is shown in Figure 10.3 below. The MC itself relies upon the services provided by the CiscoWorks Common Services software. The Common Services component provides a common environment for all of the MCs. Some of these services include data storage and management, session management, a web interface and user authentication and permission management. Before installing the Cisco IPS Management

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:35

Center, it is important to understand related software that may be prerequisites for successful installation.

Figure A.16 MC for IPS Sensors Architecture

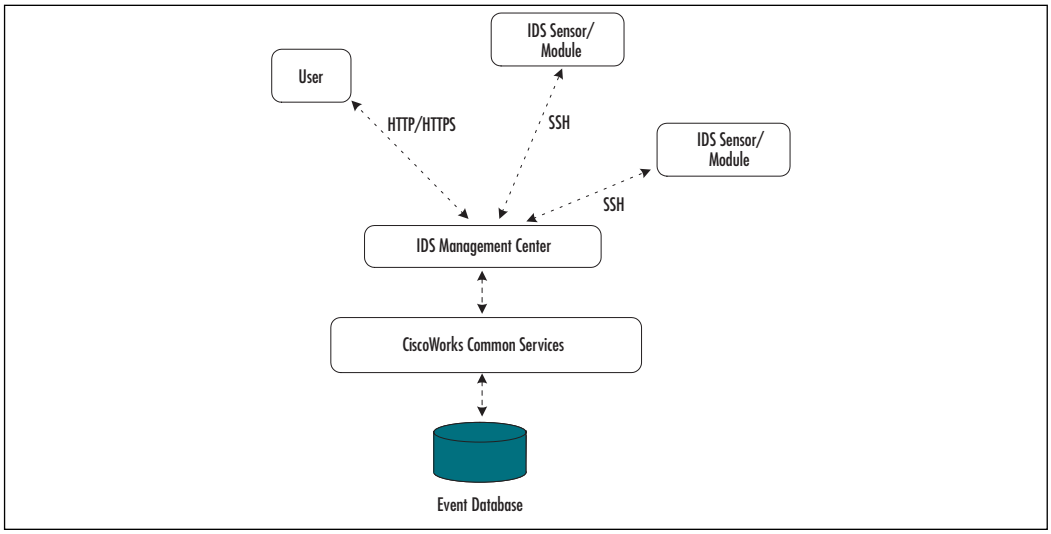

The MC for IPS Sensors is provides a web-based interface for managing and configuring Cisco IPS sensor appliances and the IPS module for the Catalyst chassis. The MC is built on top of the CiscoWorks framework allowing it to leverage the ability to define user roles. These roles provide for the definition of user management privileges including the ability to generate as well as deploy IPS configurations. The MC for IPS Sensors requires the CiscoWorks Common Services component to provide the necessary base components, software libraries and other software packages. The CiscoWorks Common Services is comprised of the following components:

- **Data Storage and Management** The Common Services data store is provided by a Sybase SQL Anytime database. Data backup, repair and restoration capabilities of the database are also provided by the Common Services package.
- **Session Management** Allows multiple users to connect to the MC and perform configuration and management tasks without data corruption or loss
- **User Management** Provides for authentication and authorization
- **Web Interface** Provided by an Apache web server allowing for connections to the MC system through a web browser. Access to the CiscoWorks server is done a secure encrypted channel over TCP port 1741. Once the user has authenticated to the CiscoWorks VMS server communication with the MC for IPS Sensors is conducted over TCP port 443.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:36

MC for IPS Sensors Installation

The MC for IPS Sensors software installs its components into the same directory as the CiscoWorks Common Services software components. This is typically in the directory: Program Files\CSCOPx. The directory structure is shown in Figure A.17 below.

Figure A.17 MC for IPS Sensors Directory Tree Structure

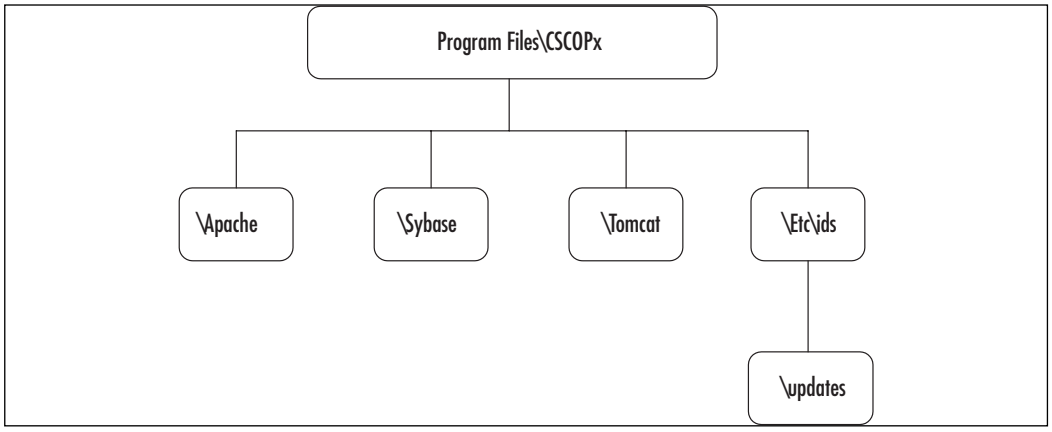

Cisco chose to use an open source program called Apache for the built in webserver for CiscoWorks. The sub-directory \Apache is where the Apache Web Server is installed and from where Apache serves the web pages that are displayed when using the MC for IPS Sensors. The Sybase sub-directory is where the Sybase SQL Anytime database is installed as well as where all data from the IPS appliances and the IDSM sensors is stored. The Tomcat is where the Tomcat application server is installed. This server provides servlets to the MC for IPS Sensors from the Common Services. The Etc\ids directory is where the MC for IPS Sensors is actually stored. The updates subdirectory is where the signature update packs are stored for the MC to push out to the sensors or to the MC itself.

MC for IPS Sensors Processes

The MC for IPS Sensors is composed of the following system processes:

- IPS_Analyzer
- IPS_Backup
- IPS_DbAdminAnalyzer
- IPS_DeployDaemon
- IPS_Notifier
- IPS_Receiver

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:37

- IPS_ReportScheduler

The IPS_Analyzer defines event rules and requests user specified notifications when appropriate. The IPS_Backup process provides for database backup and restore capabilities to the MC. The DbAdminAnalyzer applies various active database rules to the current state of the server. The IPS_DeployDemon provides for the deployment of configurations to IPS sensors. IPS_Notifier retrieves and performs MC subsystem notification requests. The IPS_Receiver receives alarms and syslog events from IPS appliance sensors and IPS modules for the Catalyst chassis and stores them in the Sybase database. As its name implies, the IPS_ReportScheduler handles the generation of reports in the MC.

VMS Component Compatibility

Most VMS components require Cisco Works Common Services to be installed in the same server. While it may seem more efficient to combine some of these VMS components on the same server, this cannot always be done for both compatibility and performance reasons.

For example, both the IPS Management Center and the Security Monitor are delivered on the same CD-ROM package. Both require CiscoWorks Common Services. The MC for IPS Sensors and the Security Monitor may be installed together or separately on different host servers. However, for optimal performance, separate installation of these two applications on different host servers is recommended.

Other VMS components that are not compatible on the same server as the IPS Management Center include the Cisco Secure Policy Manager (CSPM). To attempt this may result in the installation of a second instance of the post office process on the host server.

Client Installation Requirements

Accessing CiscoWorks VMS and IPS Management Center is accomplished through a web interface. This allows clients to access the IPS Management Center by using a browser. The minimum system requirements for clients are specified in Table A.2 below.

Table A.2 Client System Requirements

Component	Minimum Requirement
Hardware 300MHz Pentium Processor 333MHz processor	IBM PC Compatible with minimum Sun Ultra 10 or Sun SPARCstation with a
Software Edition with Service Pack 4	Windows 2000 Server, or Professional
Windows XP Professional, Service Pack 2 with Microsoft Virtual Machine	Solaris 8
Memory	256MB
Virtual Memory	400MB 512MB

Continued

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:38

Table A.2 continued Client System Requirements

Component	Minimum Requirement
Browser	Microsoft Internet Explorer 6.0, Service Pack 1 for Windows operating systems with Microsoft Virtual Machine. Netscape Navigator 7.1 on any of the following: Windows 2000 Server Professional Edition with Service Pack 4 Windows XP Professional, Service Pack 2 Netscape Navigator 4.76 for Solaris

Installation Steps

Once the prerequisite components have been verified the basic installation steps for the MC for IPS Sensors are as follows:

1. Log in as the local administrator.
2. Insert the CDROM containing the "Monitoring Center for Security and Management Center for IPS Sensors." If the installation program does not start, select **Run** from the **Start** button. Browse for the **setup** program on the CDROM drive. **Open** the **setup** program and click **OK**. If the installation program does start, click **Install** on the Installer page. Click **Next**.
3. The "Software License Agreement" page appears. Be sure that to understand the Agreement. Click **Yes** to accept the terms of the Agreement.
4. The installation now begins. To install both the MC for IPS Sensors and the Security Monitor, click the **Typical Installation** radio button. To install only the MC for IPS Sensors or the Security Monitor, click the **Custom Installation** button, and select either the **MC for IPS Sensors only** radio button or the **Security Monitor only** radio button. Click **Next**.
5. The "System Requirements" page appears. Verify that the system meets the minimum disk space and memory requirements. Click **Next**.
6. The "Verification" page appears. Verify the selected components. Click **Next**.
7. The "Select Database Location" page appears. By default, the IPS database is located in the directory where Cisco Works Common Services is installed. To specify a different directory for the database, enter a directory path in the **Database File Location** field provided. Click **Next**.
8. The "Select Database Password" page appears. Enter the database password in both the **Password** and **Confirm Password** fields. Click **Next**.
9. Either the "Select CW2000 Syslog Port" page or the "Restart" page appears.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:39

- If the Security Monitor is installed, the "Select CW2000 Syslog Port" page appears. Specify the UDP port to be used by CiscoWorks. The default value of 52514 is recommended. Click **Next**. The "Configure Communications Properties" page appears. Enter the host ID, organization ID, IP address, hostname and organization name into the appropriate fields. Click **Next**.
- If only the MC for IPS Sensors is installed, the "Restart' page appears. On the "Restart" page, select **Yes** to restart the computer. Select **No** to restart the computer at a later. Select **Finish**. The computer must be restarted before it is possible to use the MC for IPS Sensors or Security Monitor.

Getting Started

Access to the MC for IPS Sensors is provided through the Apache web server on the CiscoWorks host. This provides for easy access through either a web browser meeting the requirements defined in Table A.2 above. The CiscoWorks Apache web server listens for incoming connections on TCP port 1741 of the CiscoWorks host. To access the CiscoWorks system enter one of the following URLs:

- **http://127.0.0.1:1741** – if the MC for IPS Sensors server is the local machine
- **http://A.B.C.D:1741** – where A.B.C.D is the IP address of the MC for IPS Sensors server.

Log into the CiscoWorks Server Desktop as shown in Figure A.18 below. The default login name and password for the CiscoWorks system are:

- Default login name: **admin**
- Default password: **admin**

After entering the login **Name** and **Password**, click **Connect**. Note, it is *highly recommended* that the password of the **admin** account be changed from the default value of **admin** immediately upon installation and configuration of CiscoWorks in order to prevent unauthorized users from gaining administrative access to the CiscoWorks software.

Figure A.18 Cisco Works Login

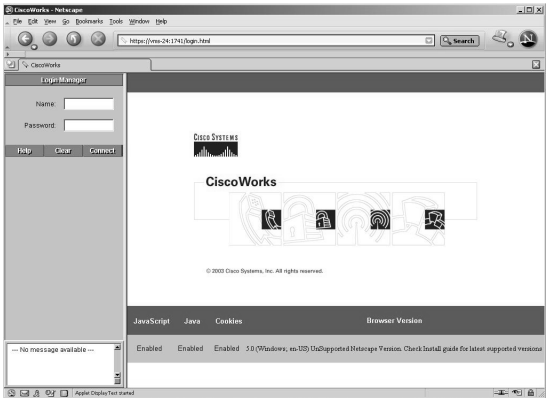

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:40

Authorization Roles

CiscoWorks provides for five different default types of accounts or authorization roles that can be created for MC for IPS Sensors users. These authorization roles and their respective privileges are summarized in Table A.3 below.

Table A.3 Authorization Roles

Authorization Role	Privileges	View	Create/Modify/Delete
Help Desk	View reports & alarms.	Cannot delete reports or alarms.	Cannot generate reports.
Approver	View reports & alarms.	Approve configurations.	Cannot delete reports or alarms. Cannot generate reports.
Network Operator	View reports & alarms.	Deploy configurations.	Delete reports and alarms. Generate reports.
Network Administrator	View reports & alarms.	Edit devices and device groups.	
System Administrator	View reports & alarms.	Edit devices and device groups.	Delete reports and alarms. Generate reports. Import lists (files) and notification scripts.

Creating accounts with different authorization roles allows an administrator to delegate different responsibilities to different IPS Management Center users. Each account holder or user can be given the authority needed to carry out his responsibilities.

Installation Verification

To verify the successful installation of CiscoWorks and the MC for IPS Sensors, select the "Server Configuration" entry on the CiscoWorks Server Desktop as shown in Figure A.19 below. Then select "About the Server" and "Applications and Versions."

Verify that the following key Cisco Works components are installed:

- **Apache**. The Apache web server is used to provide the web interface used by a client to access the IDS Management Center.
- **CWCS SQL Components.** A Sybase SQL server is used to provide the database services required by the IPS Management Center.
- **Cisco Works Common Services (CWCS).** There are a multitude of services provided by CWCS that are required by the IPS Management Center.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:41

Figure A.19 Server Configuration

CiscoWorks - Netscape
File Edit View Go Bookmarks Tools Window Help
https://vms-24:1741/login.html
Search
CiscoWorks
Logout
Help
Home
Server Configuration
About the Server
Applications and Versions
Product Overview
Administration
Diagnostics
Setup
VPN/Security Management Solution
Management Connection
Device Manager
--- No message available ---
Applet DisplayText started
Name
Version
Install Date
Installed Patches
State
Apache 1.3.27 7-18-2003 18:49:50 none ENABLED
Auto Update Server 1.1 7-21-2003 13:55:21 none ENABLED
Client Application Manager 3.0 7-18-2003 18:49:50 none ENABLED
CWCS SQL Components 7.1.3 7-18-2003 18:49:50 none ENABLED
CiscoWorks Common Services 2.2 7-18-2003 18:49:50 none ENABLED
Cisco Common Services Help 1.1 7-18-2003 18:49:50 none ENABLED
CWCS Foundation 2.2 7-18-2003 18:49:50 none ENABLED
CWCS java2 engine 1.2 7-18-2003 18:49:50 none ENABLED
CWCS Web Desktop 2.2 7-18-2003 18:49:50 none ENABLED
CWCS Utilities 1.1 7-18-2003 18:49:50 none ENABLED
Management Center for Cisco Security Agents 4.0 7-30-2003 18:31:21 none ENABLED
CiscoView 5.5 7-18-2003 18:49:50 none ENABLED
Database package 4.2 7-18-2003 18:49:50 none ENABLED
CiscoWorks Process Management package 3.5 7-18-2003 18:49:50 none ENABLED
CWCS (includes CiscoView) Help 2.2 7-18-2003 18:49:50 none ENABLED
CWCS Event Distribution System 3.2 7-18-2003 18:49:50 none ENABLED
Event Services Software 2.0 7-18-2003 18:49:50 none ENABLED

Adding Users to CiscoWorks

Adding users to the CiscoWorks system is straightforward. To add a new user to the CiscoWorks system:

1. Open the "**Server Configuration**" tab in the right side panel of the CiscoWorks interface.
2. Select "**Setup**" and then "**Security**" as shown in Figure A.19 below.
3. Select the "**Add Users**" option.
4. Enter values for the setting listed below in Table A.4 and shown in Figure A.20

Table A.4 CiscoWorks "Add Users" Information

CiscoWorks "Add Users" Setting	Information
User Name	Name of new user account to add
Local Password	Account password
Confirm Password	Password confirmation
Email	User's e-mail address (optional)
CCO Login	User's CCO login account name (optional)
CCO Password	User's CCO login account password (optional)
Confirm Password	CCO Password confirmation (optional)

Continued

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:42

Table A.4 continued CiscoWorks "Add Users" Information

CiscoWorks "Add Users" Setting	Information
Proxy Login	User's proxy server login name (optional)
Proxy Password	User's proxy server password (optional)
Confirm Password	Proxy password confirmation (optional)

Figure A.19 CiscoWorks Server Configuration>Add User

CiscoWorks - Netscape
Logout
Help
Home
Server Configuration
About the Server
Applications and Versions
Product Overview
Administration
Diagnostics
Setup
VPN/Security Management Solution
Management Connection
Device Manager
--- No message available ---
Name
Version
Install Date
Installed Patches
State
Apache 1.3.27 7-18-2003 18:49:50 none ENABLED
Auto Update Server 1.1 7-21-2003 13:55:21 none ENABLED
Client Application Manager 3.0 7-18-2003 18:49:50 none ENABLED
CWCS SQL Components 7.1.3 7-18-2003 18:49:50 none ENABLED
CiscoWorks Common Services 2.2 7-18-2003 18:49:50 none ENABLED
Cisco Common Services Help 1.1 7-18-2003 18:49:50 none ENABLED
CWCS Foundation 2.2 7-18-2003 18:49:50 none ENABLED
CWCS java2 engine 1.2 7-18-2003 18:49:50 none ENABLED
CWCS Web Desktop 2.2 7-18-2003 18:49:50 none ENABLED
CWCS Utilities 1.1 7-18-2003 18:49:50 none ENABLED
Management Center for Cisco Security Agents 4.0 7-30-2003 18:31:21 none ENABLED
CiscoView 5.5 7-18-2003 18:49:50 none ENABLED
Database package 4.2 7-18-2003 18:49:50 none ENABLED
CiscoWorks Process Management package 3.5 7-18-2003 18:49:50 none ENABLED
CWCS (includes CiscoView) Help 2.2 7-18-2003 18:49:50 none ENABLED
CWCS Event Distribution System 3.2 7-18-2003 18:49:50 none ENABLED
Event Services Software 2.0 7-18-2003 18:49:50 none ENABLED

Figure A.20 CiscoWorks Add User Web Page

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:43

The MC for IPS Sensors

If the MC for IPS Sensors installation is successful, an entry for the "Management Center" will appear. Selecting the Management Center entry will result in the appearance of the "IPS Sensors" entry. Selecting the IPS Sensors entry will bring up the "Management Center for IPS Sensors" interface as shown in Figure A.21.

Figure A.21 Management Center for IPS Sensors Page

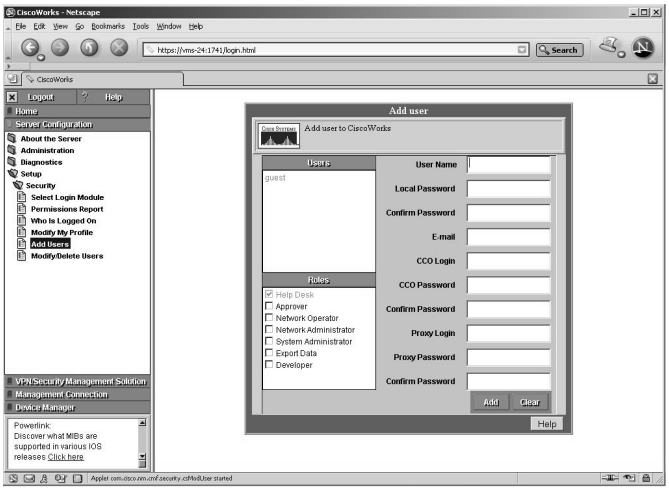

The Devices tab of this page allows for the definition of sensor groups as well as the addition or deletion of sensors from the system as described in the next section.

Setting up sensors and sensor groups

Sensors are the "eyes and ears" of the Cisco IPS Management Center. They are placed strategically at the perimeter of the network and near key resources within the enterprise. Each of the sensors deployed in the network have been configured with a unique IP address. The MC for IPS Sensors uses this IP address to communicate with the sensor. Once these sensors are deployed and assigned IP addresses they can be configured and managed from within the MC.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:44

Configuring & Implementing...

Monitoring Connections

A sensor is commonly placed on a connection to monitor traffic between the network to be protected and other networks. In Figure A.22, a protected enterprise network is comprised of two intranets: the E-Commerce network and the R&D network. Here, sensors have been deployed to monitor four different types of connections. Starting from the upper left, the sensors offer the following protection:

- **Perimeter Protection** The most common deployment for a sensor is to be placed between the network to be protected and the Internet. This is known as perimeter protection.
- **Remote Access Protection** A dial-in server is used only by employees but may still be vulnerable to external attack. A sensor is placed on the interior side of the dial-in server connection.
- **Intranet Protection** While the R&D network is an internal network, it may require a different level of security and hence a sensor is deployed between the two intranets.
- **Extranet Protection** A business partner may have similar network security policies but the level of protection may differ. A sensor is deployed between the two extranets.

Figure A.22 Monitoring Connections

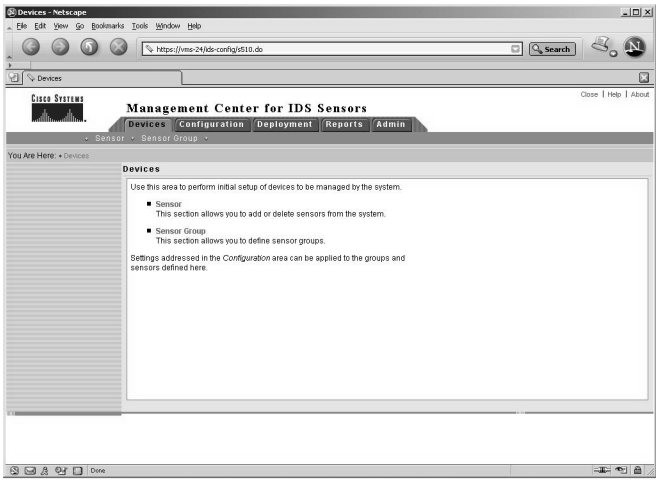
Devices - Netscape
File Edit View Go Bookmarks Tools Window Help
Search
Devices
Cisco Systems
Management Center for IDS Sensors
Devices
Configuration
Deployment
Reports
Admin
Sensor
Sensor Group
You Are Here: • Devices
Devices
Use this area to perform initial setup of devices to be managed by the system.
Sensor
This section allows you to add or delete sensors from the system.
Sensor Group
This section allows you to define sensor groups.
Settings addressed in the Configuration area can be applied to the groups and sensors defined here.
Done

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:45

MC for IPS Sensors Hierarchy

The MC for IPS Sensors maintains a hierarchy of sensors, sensor groups and sensor subgroups. Groups provide the capability of managing multiple sensors performing similar functions. Rather than configuring each sensor individually, the MC for IPS Sensors allows for the configuration of groups of sensors. This dramatically reduces the administrative burden on security personnel. Figure A.23 illustrates an example of an MC for IPS Sensors sensor group hierarchy. At the top of the group hierarchy is the "Global" group. There can be many levels of groups and sensors under the Global group. Each of the lower level groups, subgroups and sensors are added manually.

Figure A.23 MC for IPS Sensors Hierarchy

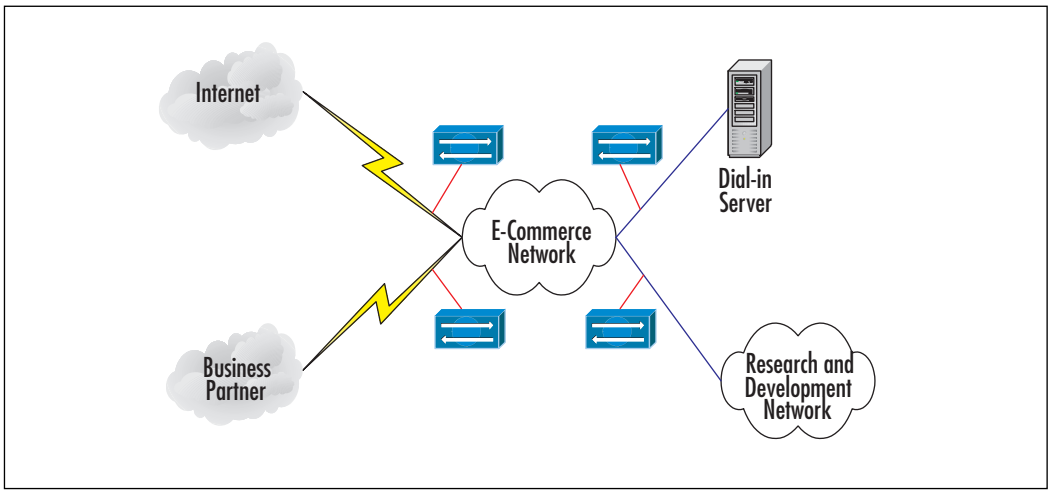

Creating Sensor Subgroups

A sensor subgroup can be added to any group including the Global group. The following steps can be used to create a sensor subgroup:

1. From the "Management Center for IPS Sensors" page (Figure 10.7), select the **Devices** tab then select **Sensor Group**. The "Sensor Group" page will appear as shown in Figure 10.11 above.
2. The Sensor Group page displays a tree of multiple levels of sensor groups and sensors. At the present, there is a Global group as well as three subgroups: Core, Internet, and VPN. Select the name of the group under which you the new subgroup will appear. Click the **Create Subgroup** button.
3. The "Add Group" page appears as shown in Figure A.24. Enter the new subgroup's name in the **Group Name** field. Describe in the **Description** field the

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:46

new group. Under settings, select the parent group's settings or copy the settings from a group in the pull-down menu.

4. Click **Ok** to create the new subgroup.

Figure A.24 Add Group Page

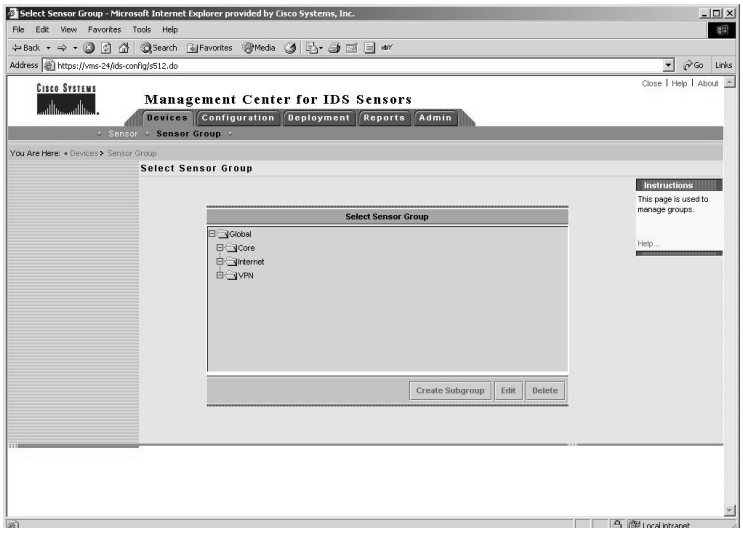

The "Sensor Group" page reappears and contains the newly created group. In Figure A.25, this new group is named **Campus**.

Figure A.25 Sensor Group Page with new Subgroup

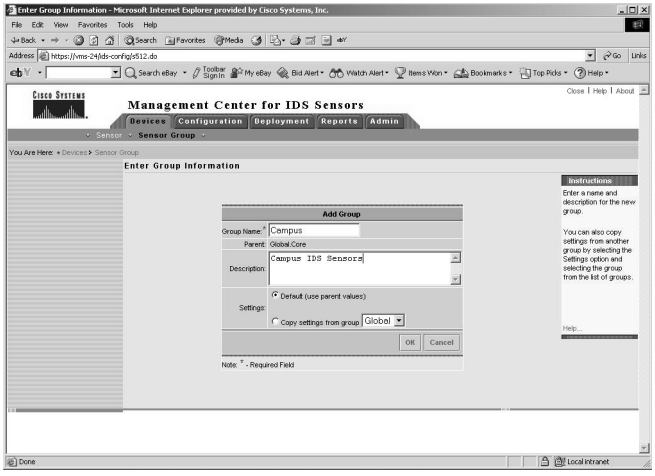

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:47

Adding Sensors to a Sensor Group

A sensor can be added to any group including the Global group. To add a sensor to the Global group or a subgroup use the following procedure:

1. From the "Management Center for IPS Sensors" page (Figure A.26) select the **Devices** tab then select **Sensors**.
2. The "Sensor" page will appear as shown in Figure A.26. Click the **Add** button.

Figure A.26 Sensor Page

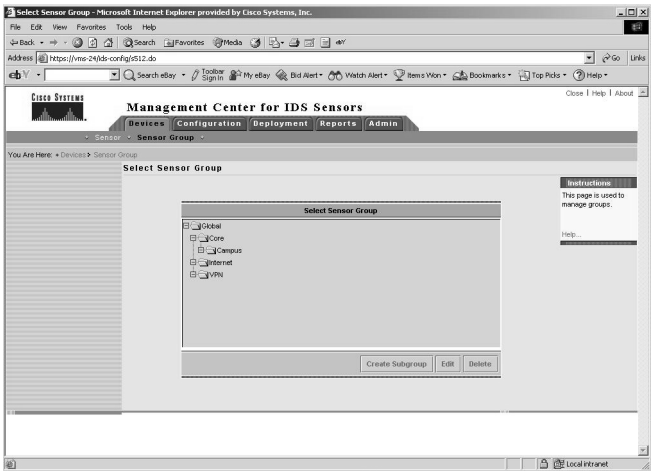

3. The "Select Group" page will appear as shown in Figure A.27. Select the Group to add the sensor to and click **Next**.
4. The "Enter Sensor Information" page appears as shown in Figure A.28 below. Enter the **IP Address** of the sensor, the **NAT Address** of the sensor if one exists, and the **Sensor Name.** To retrieve sensor settings directly from the sensor, select the **Discover Settings** check box. Enter the **User ID** and **Password** for Secure Shell (SSH) communications. For sensor appliances and IPS modules, the default user ID is **cisco**. The default password for the account is **cisco**. It is also possible to authenticate to the IPS sensor using an SSH public/private key pair. To use existing SSH keys, check the **Use Existing SSH keys** checkbox. However, do not select this option if the sensor is to be used as a master blocking sensor. Once the information has been entered click **Next** to move on to the final step.

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:48

Figure A.27 Select Group Page

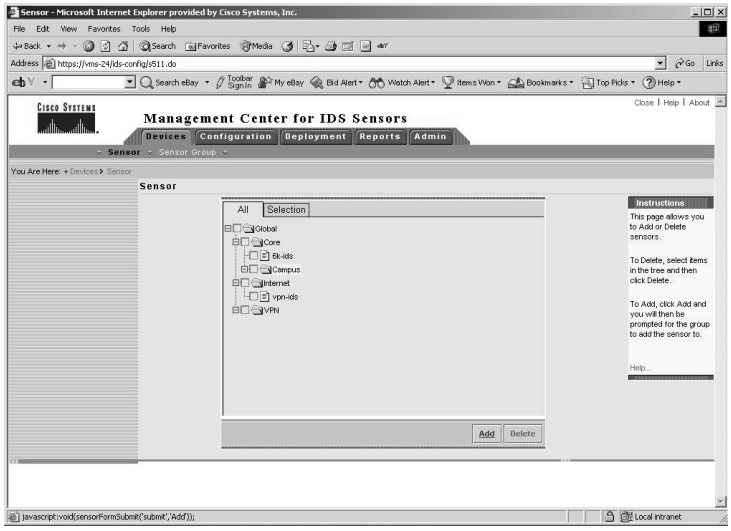

Figure A.28 Enter Sensor Information Page

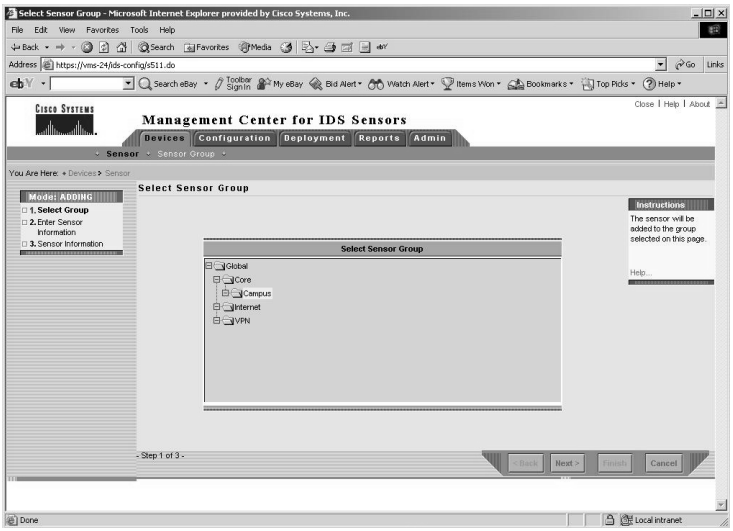

5. The "Sensor Information" page appears as shown in Figures A.29 and A.30. From the **Version** pull-down menu, select the sensor software version installed on the sensor. Enter a text **Comment**. For sensors running the IPS sensor software version 3.x additional information needs to be entered. This information includes the sensor **Host ID** which is typically the last octet of the sensor's IP address. Enter the **Org Name** using only lower case letters. Enter the **Org ID**. The default is

403_Ent_DMZ_AA.qxd 10/25/06 11:22 AM Page A:49

100. Within a **Postoffice** domain, no sensor or sensor group, the Org ID/Host ID pair must be unique. For Sensor software version 4.x and later a text comment need only be entered in the **Comment** field. Click **Finish**.

Figure A.29 Sensor Information Page for Sensor OS Version 3.x

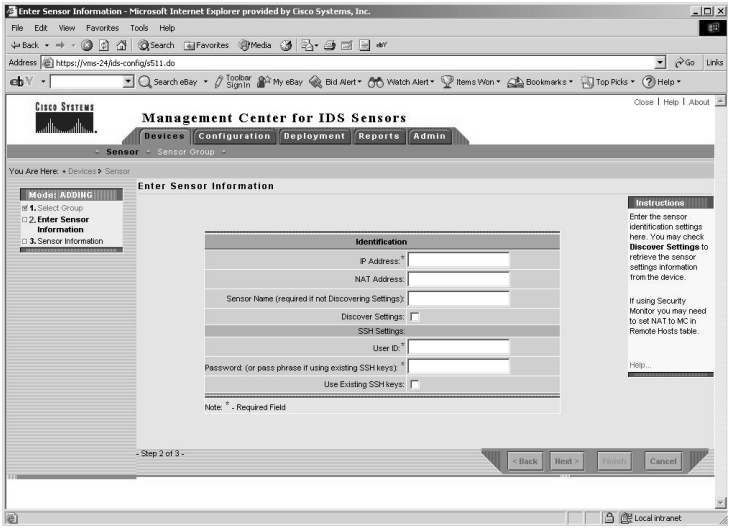

Figure A.30 Sensor Information Page for Sensor OS Version 4.x

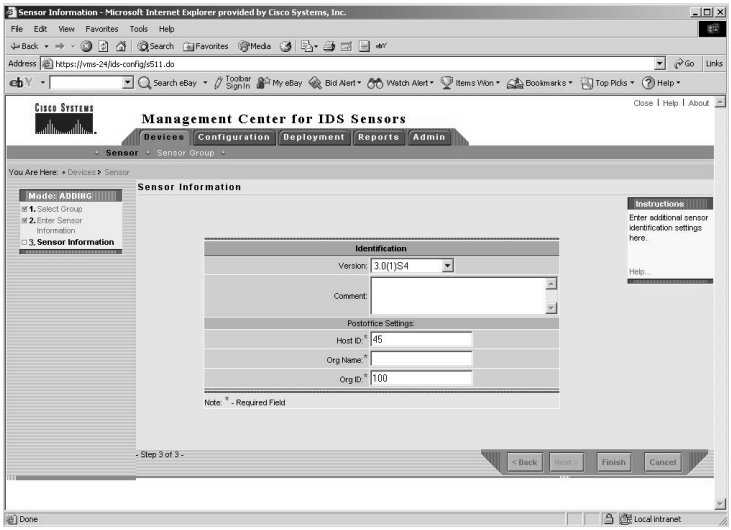

6. The "Sensor" page reappears, updated with an entry for the new sensor you have added as shown in Figure A.31.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:50

Figure A.31 Updated Sensor Page

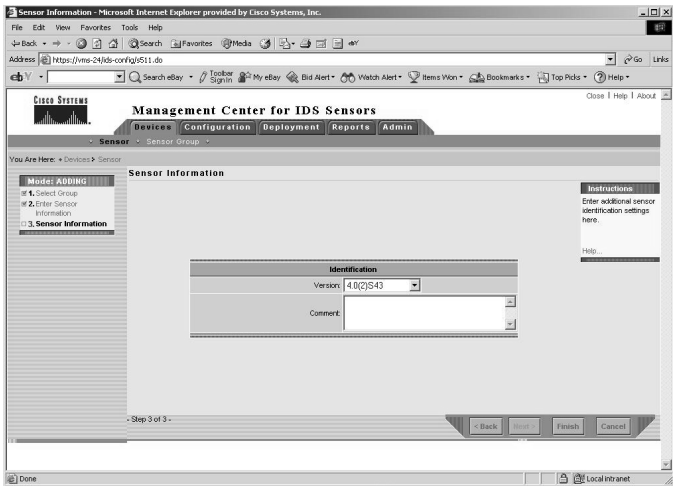

Deleting Sensors from a Sensor Group

A sensor can be deleted from any group including the Global group. Use the following steps to delete a sensor from a subgroup:

1. From the "Management Center for IPS Sensors" page (Figure A.19), select the **Devices** tab and select **Sensors**.
2. The "Sensor" page will appear as shown in Figure A.32. Check the box in front of the entry for the sensor to delete. In this case, the sensor to be deleted is call "thorin." Click the **Delete** button.

Figure A.32 Sensor Page

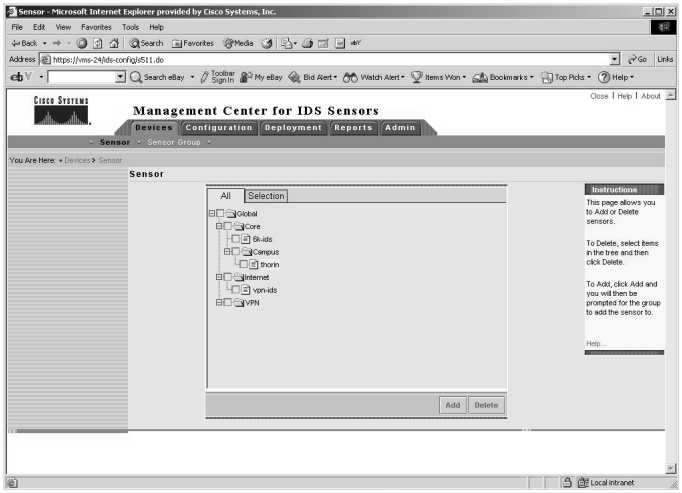

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:51

3. The "Sensor" tree page appears as shown in Figure A.33. Note that the sensor, known as "thorin" has been removed from the tree.

Figure A.33 Sensor Tree Page

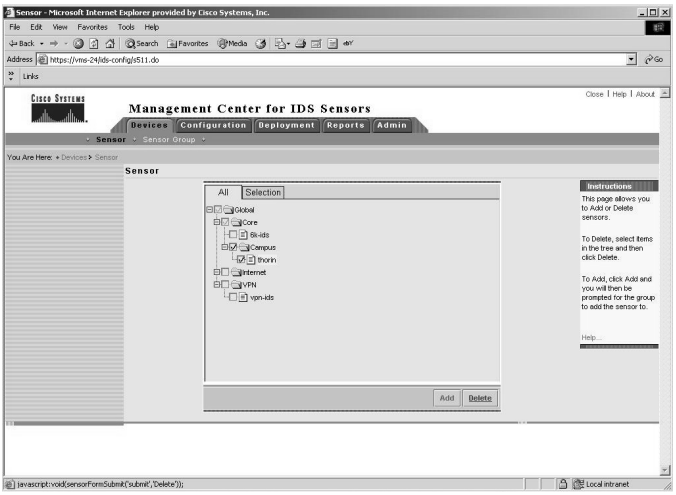

Deleting Sensor Subgroups

As with sensors, sensor subgroups can be deleted from any group including the Global group. Use the following steps to delete a sensor subgroup:

1. From the "Management Center for IPS Sensors" page (Figure A.19), select the **Devices** tab and select **Sensor Group**.
2. The "Sensor Group" page will appear as shown in Figure A.34. In the tree, select the subgroup to delete and click the **Delete** button.

Figure A.34 Sensor Group Page

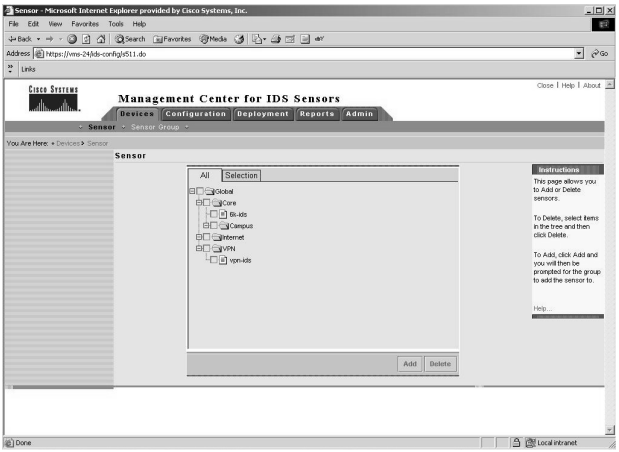

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:52

Configuring Signatures and alarms

Network intrusions are scans, attacks upon or misuses of the network resources. To detect network intrusion, the Cisco IPS Sensors use a signature-based technology. Every network attack has an order or a pattern to the bytes in the traffic stream between the attacking system and the target. These bytes represent a "fingerprint" or "signature" of the attack. By comparing the pattern of bytes in a given traffic stream between two hosts against a database containing various known signatures for network attacks the IPS is able to determine when an attack has occurred. Each signature specifies the type of attack the sensor detects and reports. As a sensor scans the network packets, the rules allow it to detect patterns that match a known attack.

The MC for IPS Sensors allows the operator to specify which signatures should be enabled. Additionally, the response action the IPS sensor initiates, whether it is simply raising an alarm on the Security Monitor console or initiating a TCP RST, is also determined based on what is specified in the signature. Tuning IPS signatures is one of the more important features of the MC for IPS Sensors. Improperly tuned IPS sensors account for the great majority of false positive alarms (alarms that are raised by the IPS in response to benign network traffic) and result in potential mistrust of the IPS system by the security personnel.

Configuring Signatures

Signatures are divided into six groups:

1. General (embedded)
2. TCP Connection
3. UDP Connection
4. String-Matching
5. Access Control List (ACL)
6. Custom.

To provide an example of how to configure and tune signatures, we will use a general signature for a configuration and tuning exercise.

Configuring General Signatures

General signatures are signatures that are embedded in the sensor software itself. IPS end-users cannot add or delete general signatures, but the end-user can enable or disable them and configure the response to attacks that fit the general signatures. The following steps can be used to configure a general signature:

1. From the "Management Center for IPS Sensors" page, select **Configuration | Settings**.
2. A "Table of Contents" page appears. Select the **Object Selector** handle.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:53

3. In the Object Selector, select the sensor containing the general signature to configure. The Object Selector will close and redisplay the Table of Contents.
4. In the Table of Contents, select **Signatures | General.** The General page will appear as shown in Figure A.35.

Figure A.35 General Signature Page

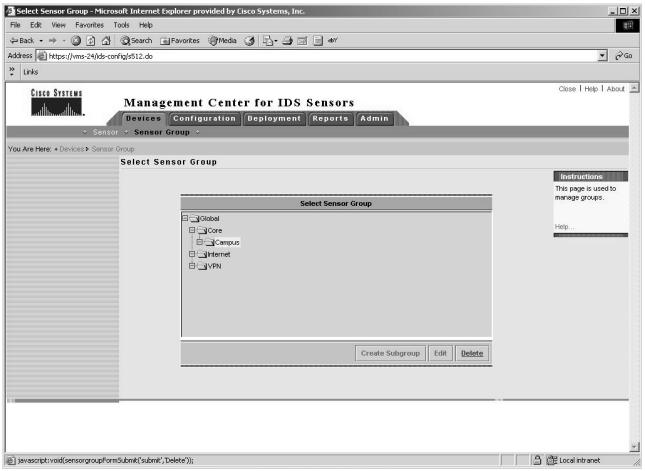

5. Click on the link for the signature group to be modified. This results in the display of the Group Signature page listing all of the signatures within the selected group as shown in Figure A.36.
6. Select the signature to configure by checking the corresponding box and clicking **Edit**.
7. The "Edit Signatures" window appears as shown in Figure A.37 and shows the name of the signature to configure. To enable or disable the signature, check or uncheck the **Enable** box.

Configuring Alarms

The severity of an alarm as well as the actions to be taken when an event matches a signature can be specified by editing the signature.

1. To change the severity of an attack that matches this signature, select a **Severity** from the pull-down menu:

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:54

Figure A.36 Group Signature Page

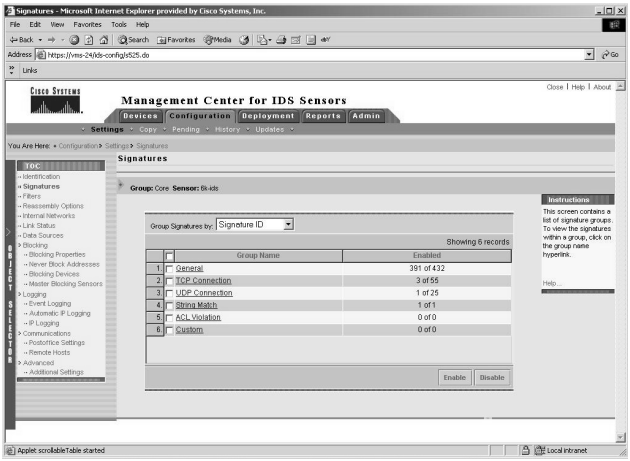

Figure A.37 Edit Signature Page

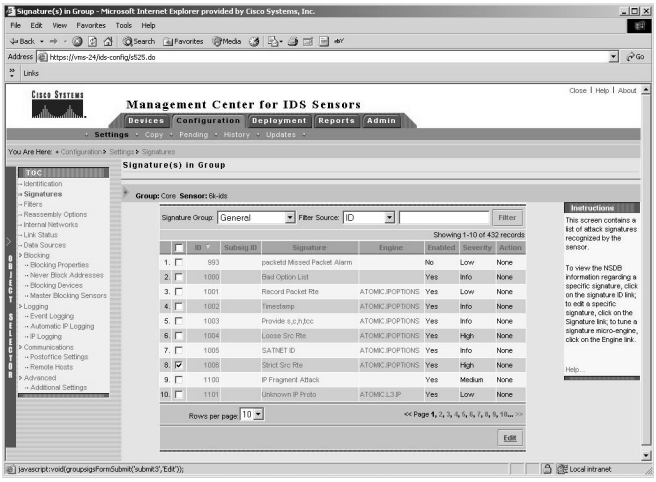

- **Info** indicates an event that results from normal activity
- **Low** indicates an attack that is mild in severity. The Security Monitor Event Viewer will display this type of attack with a green icon.
- **Medium** indicates an attack that is moderately severe. The Security Monitor Event Viewer will display this type of attack with a yellow icon.
- **High** indicates an attack that is highly severe. The Security Monitor Event Viewer will display this type of attack with a red icon.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:55

- Note the options to the right of the **Actions** label. Depending on the signature, you may specify one or more of the following actions to be taken when a signature matches an event:
- **Log** stands for "IP Log" and will generate an IP session log with information about the attack.
- **Reset** stands for "TCP Reset" and will reset the TCP session in which the attack signature was detected.
- **Block** will cause the sensor to issue a command to a PIX firewall or Cisco router. That firewall or router will block packets from the attacking host or network and keep those packets from entering the protected network.

Tuning General Signatures

Signatures are tuned to minimize false alarms or "false positives." False positives are alarm indicators of an attack when either benign or standard activity is present. A false positive may result from normal network activity in which a network management station polls or scans network devices to ascertain their status. This polling activity is similar to the scanning employed by hackers against a targeted network. Additionally, a false positive may occur when an attacker attempts to use an exploit against a host whose software is not vulnerable to the exploit (i.e. using a Microsoft IIS exploit against an Apache web server).

To tune a signature, return to the "General" signature page shown in Figure 10.21. For the signature to be tuned, select the signature link in the **Engine** column of the table. This will bring up the **Tune Signature** page as shown in Figure A.38 below.

Figure A.38 Tune Signature Page

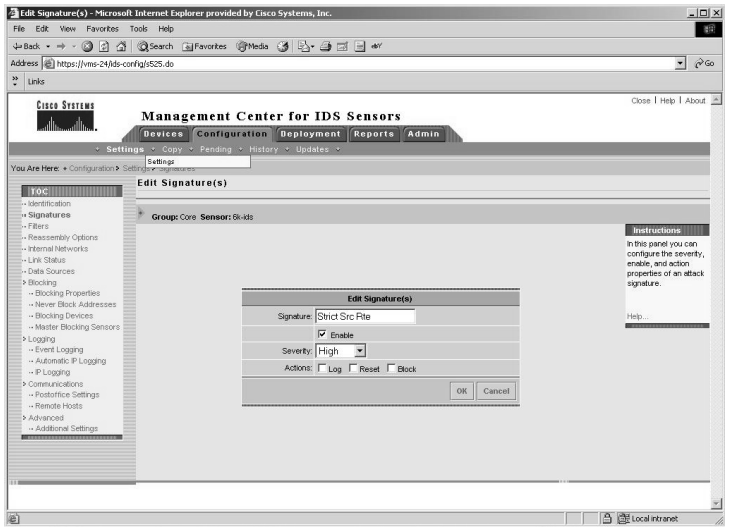

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:56

There are three columns in the **Tune Signature Parameters** table: **Parameter Name**, **Value**, and **Default**. Each one can be modified to an appropriate, desired value. Use the following procedure to tune a given parameter in a procedure:

1. Select the radio button for the parameter to be tune in the **Parameter Name** column them select **Edit** as shown in Figure A.39 below

Figure A.39 Signature Parameters Page

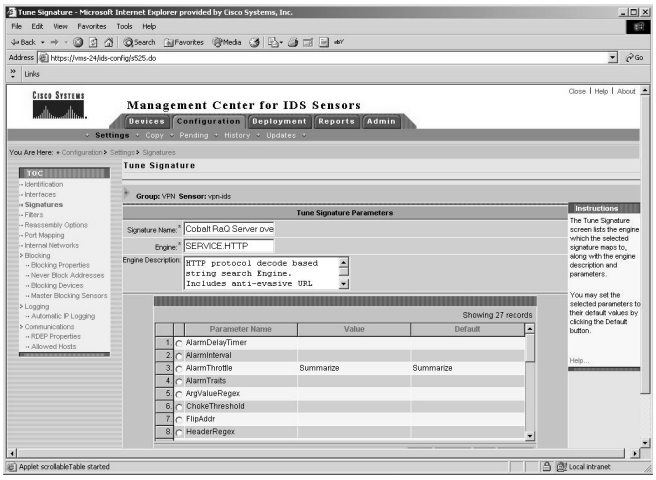

2. Enter a value for the parameter in the **Value** field as shown in Figure A.40 below.
3. Enter an optional description for the signature parameter in the **Description** field.

Figure A.40 Signature Parameters Page

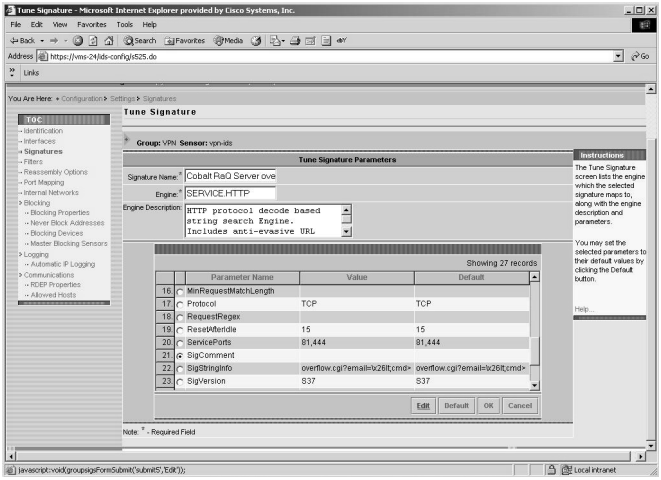

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:57

4. To accept the changes, click the **OK** button. The "Tune Signature" page will redisplay.
5. On the Tune Signature page, click **OK** to accept the changes. The General page will reappear.

How to Generate, Approve and Deploy IPS Sensor Configuration Files

The previous section, "Configuring Signatures and Alarms," covered how to select the proper values for the sensor settings and signature settings. The next step in using the MC for IPS Sensors is to review and generate the configuration files that contain those settings. Once the configuration files for the IPS sensors have been generated they need to be reviewed by the appropriate personnel and then deployed to the sensors. This section, cover how to review and generate the IPS sensor configuration files as well as how to approve and deploy the configuration files to the sensors.

Reviewing Configuration Files

Changes to file settings are placed in a pending status before they are committed to the IPS Database. The following steps can be used to review the pending changes and commit it them to the database:

1. From the "Management Center of IPS Sensors" page in Figure A.19, select **Configuration | Pending**. The "Pending" configurations page appears as shown in Figure A.41.

Figure A.41 Pending Configurations Page

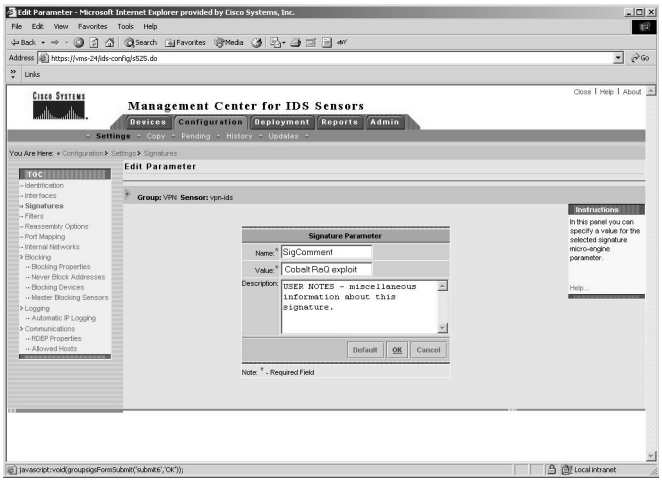

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:58

2. Check the box associated with the sensor whose configuration is to be saved in the IPS Database.
3. Click **Save** to save the configuration in the IPS Database or click **Delete** to delete it.

Generating Configuration Files

To generate a configuration file is to take a file of sensor configuration settings that is stored in the IPS Database and prepare it for deployment to the sensor itself. Generating a configuration file starts with the "Management Center for IPS Sensors" page as shown in Figure A.19.

1. From the "Management Center of IPS Sensors" page shown in Figure A.19, select **Deployment | Generate**.
2. The "Generate" page appears as shown in Figure A.42. To generate a configuration file for a specific sensor, select that sensor from the tree and click **Generate**. Once the configuration file has been generated it is now ready for the approval process.

Figure A.42 Generate Page

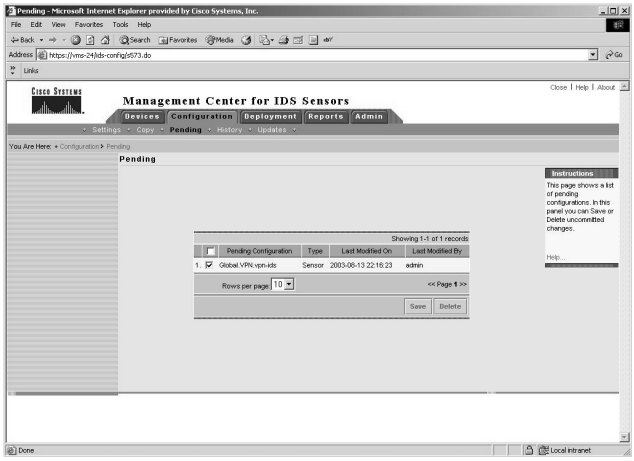

Approving Configuration Files

CiscoWorks allows for a separation of duties among user roles. This makes it possible to assign the approval of configuration files and other actions to a specific account. By separating various functions among different accounts CiscoWorks allows for a "checks-and-balance" system whereby administrators are able to verify configurations for network equipment. This is especially important in IPS because an error in the configuration file for an IPS sensor may result in the sensor not identifying an attack.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:59

1. From the "Management Center for IPS Sensors" page in Figure A.19, select **Deployment | Approve**.
2. The "Approve" page appears as shown in Figure A.43. To approve the configuration generated, check the corresponding box and click the **Approve** button.

Figure A.43 Approve Page

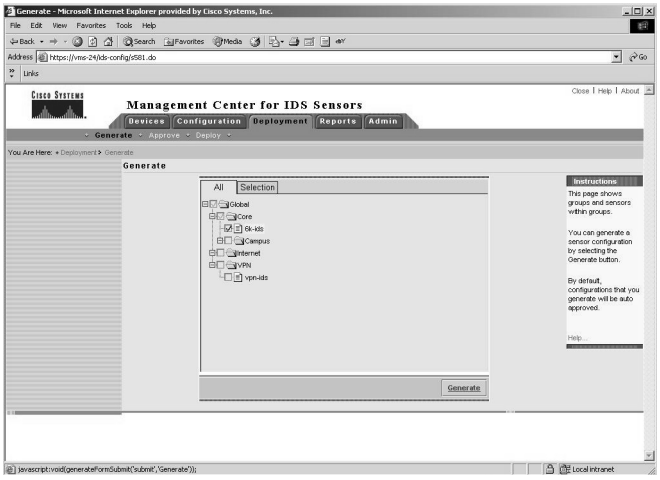

3. To view a selected IPS configuration file before approving it check the corresponding box to the right of the configuration file name and click the **View** button.
4. To delete an IPS configuration without approving it check the corresponding box to the right of the configuration file name and select the **Delete** button.

Deploying Configuration Files

To deploy a configuration file is to send an approved file of sensor configuration settings from the IPS Database to the sensor itself. Use the following steps to deploy a configuration file:

1. From the "Management Center for IPS Sensors" page select **Deployment | Deploy.** From the "Table of Contents," then select **Submit.**
2. The "Submit" page appears as shown in Figure A.44. From the tree, check the box next to the sensor name that the configuration file is to be deployed.
3. The "Select Configuration" page appears. Select a sensor configuration by checking the corresponding box and click **Next.**

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:60

Figure A.44 Submit Page

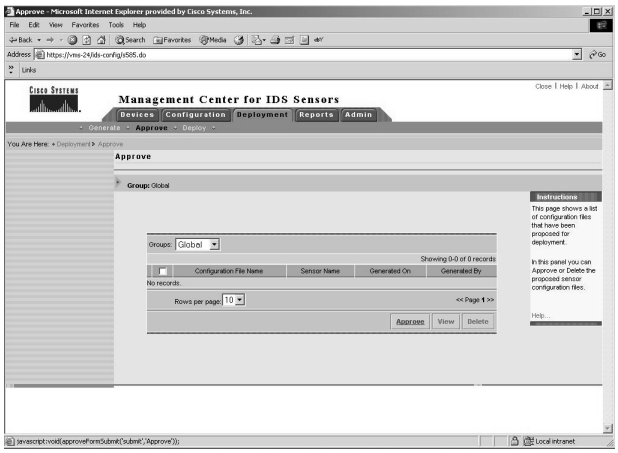

4. The "Enter Job Properties" page appears. Under **Schedule Type,** enter the name of the job from the **Job Name** field.
5. The job will deploy the configuration to the selected sensor. To start the job immediately, click the **Immediate** button. To schedule the job to execute at a later time, click the **Scheduled** radio button and select the desired options.
6. Click the **Finish** button.
7. The "Submit" page appears. To verify the scheduled job return to the "Management Center for IPS Sensors" page as shown in Figure A.19. Select **Deployment | Deploy.** From the "Table of Contents," select **Pending.** The "Pending" jobs page appears as shown in Figure A.45. On this page, it is possible to edit a pending deployment or delete it by using the Edit and Delete buttons.

Figure A.45 Pending Jobs Page

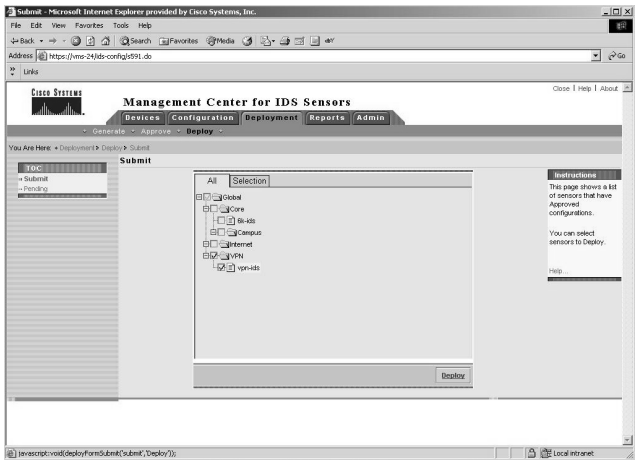

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:61

Configuring Reports

Reports provide for a summarization of the various activity and configuration of the deployed IPS sensors as well as the IPS Management Center itself. This is crucial when managing and monitoring an enterprise-wide deployment of IPS since it becomes impractical to query each IPS sensor manually in order to determine its status. The IPS Management Center can produce reports, known as audit reports, which provide information about network configuration activities managed with the Cisco MC for IPS Sensors. These reports can be generated from the **Reports** tab of the "Management Center for IPS Sensors" page as shown in Figure 10.6.

Additional reports are available from the Security Monitor. The Security Monitor is a closely related but separate product that receives real time communications from the sensors. When the IPS Management Center and the Security Monitor are installed in the same host system, the audit report templates are shared between the two products.

Audit Log Reports

There are six types of audit reports that are available from the IPS Management Center.

- Subsystem Report
- Sensor Version Import Report
- Sensor Configuration Import Report
- Sensor Configuration Deployment Report
- Console Notification Report
- Audit Log Report

The following sections examine each report in detail.

Subsystem Report

The Cisco Intrusion Detection System has many subsystems. These subsystems include the Management Center, the Security Monitor, and other subsystems. The Subsystem Report shows audit records separated and ordered by subsystem. The entries in the Subsystem Report can be filtered by Event Severity, Date/Time, and Subsystem.

Sensor Version Import Report

The IPS Management Center tracks the version identifier of each sensor. When the version identifier of a sensor is imported to the MC for IPS Sensors, an audit record is generated. The audit record indicates the success or failure of the import operation. The entries in the Sensor Version Import Report can be filtered by Device, Event Severity, and Date/Time.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:62

Sensor Configuration Import Report

IPS sensor configurations are often imported into the IPS Management Center for viewing or editing. Audit records are generated when this import operation is executed. The audit record indicates the success or failure of the import operation. The entries in the Sensor Configuration Import Report can be filtered by Device, Event Severity, and Date/Time.

Sensor Configuration Deployment Report

File configurations containing new settings are often deployed to the sensors. Audit records are generated when this deployment operation is executed. These records can indicate successful deployment or provide error messages. The entries in the Sensor Configuration Deployment Report can be filtered by Device, Event Severity, and Date/Time.

Console Notification Report

The IPS Notification subsystem generates console notification audit records. The entries in the Console Notification Report can be filtered by Event Severity and Date/Time.

Audit Log Report

The Audit Log Report displays audit records by the IPS server and by the IPS application. This report template provides a broad, non-task-specific view of audit records in the database. The entries in the Audit Log Report can be filtered by Task Type, Event Severity, Date/Time, Subsystem and Application.

Generating Reports

Reports can be generated immediately or scheduled at a later time. We can generate a report by starting from the "IPS Management Center for IPS Sensors" page and selecting the "Reports" tab. The resulting page is shown in Figure A.46.

Figure A.46 Management Center for IPS Sensors Page

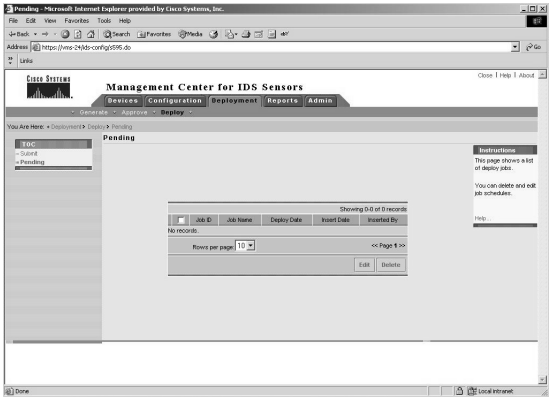

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:63

To generate a report, follow these steps:

1. From the "Reports" page, select **Generate**.
2. The "Select Report" page appears. Select the type of report that to generate and click **Select.**
3. The "Report Filtering" page appears. Enter the report parameters for the report selected and click **Next.**
4. The "Schedule Report" page appears. In the **Report Title** field, specify a name for the report. Select a radio button to schedule the report:
   - **Run Now** will generate the report immediately.
   - **Schedule for Later** will allow the specification of when the report will be generated including the generation of reports on regular intervals.
5. The Email report to field allows the specification of an email address of a report recipient. Click Finish.
6. To view the reports scheduled for generation, from the "Management Center for IPS Sensors" page, select **Reports | Scheduled.**

Viewing Reports

To view a generated report, start from the "Management Center for IPS Sensors" page and follow the following steps:

1. Select **Reports | View**.
2. The "Choose Completed Report" page appears. Check the box corresponding to the title of the report to view and click **View**.

Exporting Reports

To export a generated report to an HTML file, start from the "Management Center for IPS Sensors" page and follow the following steps:

1. Select **Reports | View**.
2. The "Choose Completed Report" page appears. Check the box corresponding to the title of the report you want to view and click **Open in Window**.
3. Depending on the browser that appears, select **File | Save As** or **Save File.** Browse to the location where the file is to be saved, enter a file name and click **Save**.

Deleting Generated Reports

To delete a generated report, start from the "Management Center for IPS Sensors" page and follow the following steps:

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:64

1. Select **Reports | View**.
2. The "Choose Completed Report" page appears. Check the boxes corresponding to the titles of the reports to delete and click **Delete**.

Edit Report Parameters

To edit the schedule for a report or the parameters for a scheduled report, start from the "Management Center for IPS Sensors" page and follow the following steps:

1. Select **Reports | Scheduled**.
2. The "Edit Scheduled Reports" page appears. Check the box corresponding to the title of the report template to edit and click **Edit**.
3. A new page appears displaying the report parameters. Change any report parameter and click **Finish**.

Example of IPS Sensor Versions Report Generation

This section details the generation of an example report. Use the following procedure to generate and view reports:

1. Select **Reports | Generate** to select the type of report to be generated from the **Select Reports** page.
2. In the **Select Reports** page, choose one of the report types desired as shown in Figure A.47 below and click **Select**.

Figure A.47 Report Selection Page

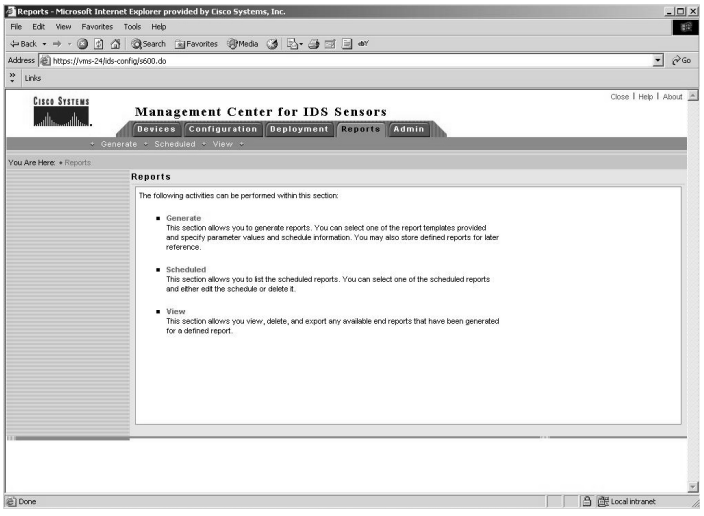
Reports - Microsoft Internet Explorer provided by Cisco Systems, Inc.
File Edit View Favorites Tools Help
Address https://vms-24/ids-config/s600.do
Links
Close | Help | About
Cisco Systems
Management Center for IDS Sensors
Devices Configuration Deployment Reports Admin
Generate Scheduled View
You Are Here: Reports
Reports
The following activities can be performed within this section:
Generate
This section allows you to generate reports. You can select one of the report templates provided and specify parameter values and schedule information. You may also store defined reports for later reference.
Scheduled
This section allows you to list the scheduled reports. You can select one of the scheduled reports and either edit the schedule or delete it.
View
This section allows you view, delete, and export any available and reports that have been generated for a defined report.
Done
Local intranet

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:65

3. The next step is to schedule the report. In the **Schedule Report** page shown in Figure A.48 the report generation can be scheduled to occur immediately, with the **Schedule Options | Run Now** option, or for some later period (**Schedule Options | Schedule for Later**).
4. Select the **Finish** button to generate the report.

Figure A.48 Report Selection Page

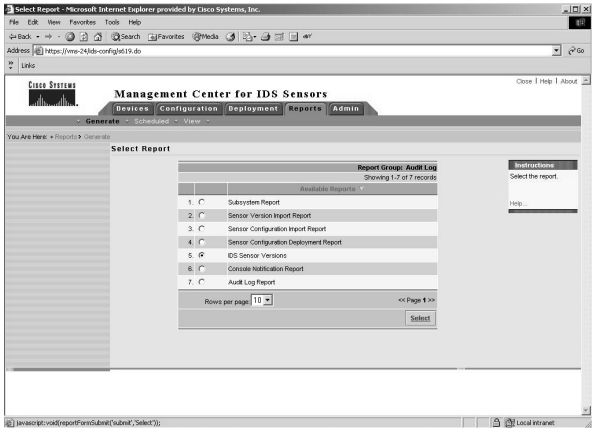

5. Once the report generation is complete the report title will appear in the list of completed reports. Select the checkbox (or checkboxes) of the report (or reports) to view and select **View** as shown in Figure A.49 below

Figure A.49 Report View Selection Page

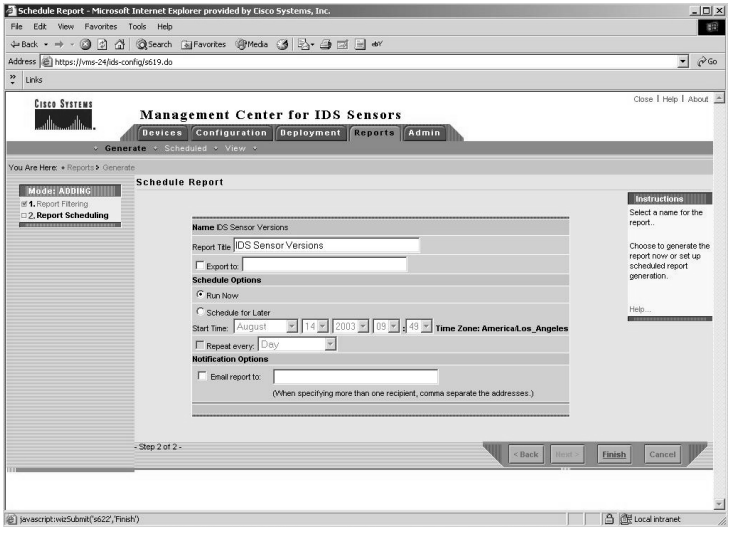

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:66

Security Monitor Reports

While the IPS Management Center can provide audit log reports, information about network activities detected by the IPS Sensors are usually provided by the Security Monitor. To access the Security Monitor from the Cisco Works Desktop, select the "Monitoring Center" and then the "Security Monitor" as shown in Figure A.50.

Figure A.50 Security Monitor

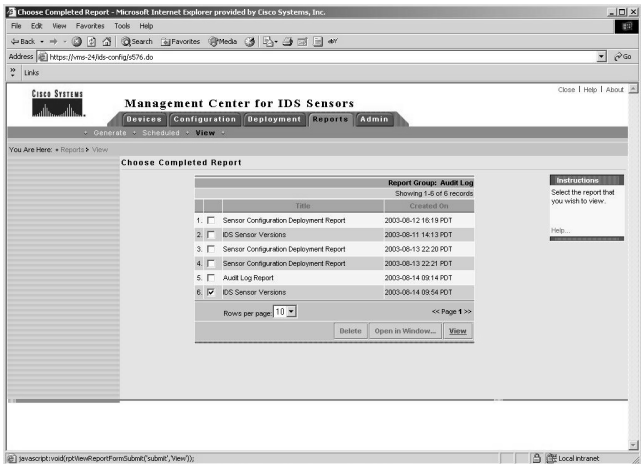
Choose Completed Report - Microsoft Internet Explorer provided by Cisco Systems, Inc.
Management Center for IDS Sensors
Devices
Configuration
Deployment
Reports
Admin
Choose Completed Report
Report Group: Audit Log
Showing 1-6 of 6 records
Instructions
Select the report that you wish to view.
1. Sensor Configuration Deployment Report 2003-08-12 16:19 PDT
2. IDS Sensor Versions 2003-08-11 14:13 PDT
3. Sensor Configuration Deployment Report 2003-08-13 22:20 PDT
4. Sensor Configuration Deployment Report 2003-08-13 22:21 PDT
5. Audit Log Report 2003-08-14 09:14 PDT
6. IDS Sensor Versions 2003-08-14 09:54 PDT
Rows per page 10
Delete
Open in Window...
View

To access reports provided by the Security Monitor, select the **Reports** tab and then the **View** entry. This will bring up the **Completed Reports** menu as shown in Figure A.51.

Figure A.51 Security Monitor Completed Reports

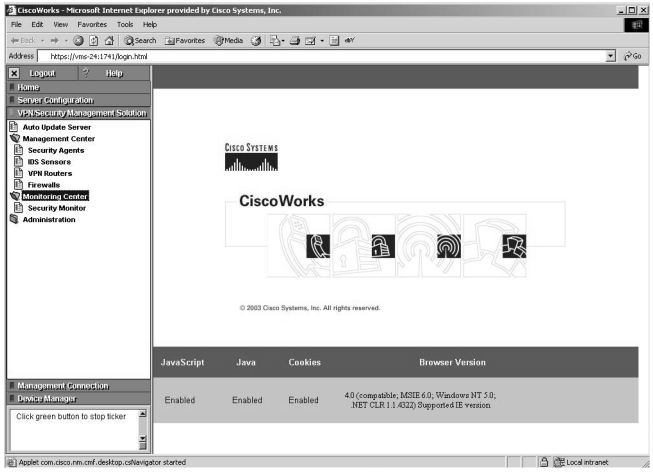

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:67

To select a report for viewing, check the box next to the report and click on the **View** button.

Administering the Cisco MC for IPS Sensors server

The administration of the Cisco MC for IPS Sensors server is comprised of tasks associated with the IPS Database and other global tasks. This encompasses:

- Operations with database rules
- Updating sensor software and signature release levels
- Defining the e-mail server settings
- Setting the configuration file approval method.

Database Rules

Database rules are used to configure the Cisco IPS Management Center to take an action at daily intervals or when a database threshold has been reached. These actions to be taken may include: sending an email notification, logging a console notification event, or executing a script.

Adding a Database Rule

To add a database rule, start from the "Management Center for IPS Sensors" page, select the "Admin" tab and "Database Rules" as shown in Figure A.52 and use the following steps:

Figure A.52 Data Base Rules

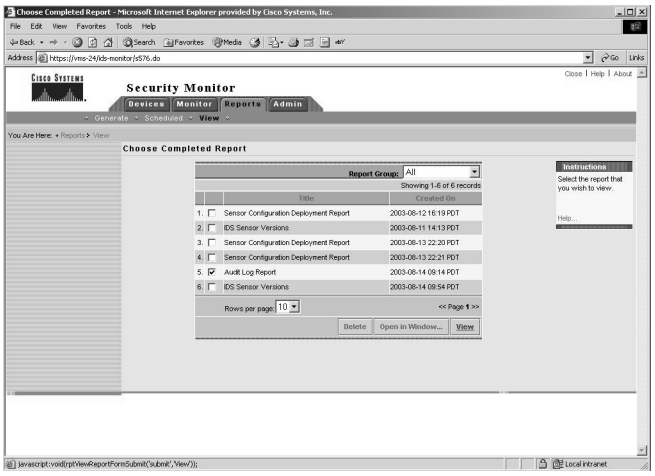

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:68

1. Select **Admin | Database**.
2. The "Database Rules" page appears. Click **Add**.
3. The "Specify the Trigger Conditions" page appears. Specify the threshold to trigger Security Monitor to take an action. The following triggers can be specified with checkboxes:
   - **Database used space greater than (megabytes)**. This will trigger an action when the database reaches a size in megabytes that is specified in the next field.
   - **Database free space less than (megabytes)**. This will trigger an action when the database free space drops to a size in megabytes that is specified in the next field.
   - **Total IPS events**. This will trigger an action when the total number of IPS events in the database reaches the number is specified in the next field.
   - **Total SYSLOG events**. This will trigger an action when the total number of SYSLOG events in the database reaches the number is specified in the next field.
   - **Total events**. This will trigger an action when the total number of events in the database reaches the number is specified in the next field.
   - **Daily beginning**. This will trigger an action to occur daily beginning on the date and time specified.

   In the **Comment** field you may enter a description of the Database Rule. Click **Next.**
4. The "Choose the Actions" page appears. More than one action can be selected via the following checkboxes:
   - Notify via Email
   - Log a Console Notification Event
   - Execute a Script.
5. Click **Finish.**

Editing a Database Rule

To edit a database rule, start from the "Management Center for IPS Sensors" page as shown in Figure 10.29 and follow these steps:

1. Select **Admin | Database**.
2. The "Database Rules" page appears. Select the radio button corresponding to the rule to edit and click **Edit**.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:69

3. The "Specify the Trigger Conditions" page appears. Select the radio button corresponding to the rule to edit and click **Edit**. Change the field that that is to be revised and click **Next**.
4. The "Choose the Actions" page appears. Make the desired changes and click **Finish**.

Viewing a Database Rule

To view a database rule, start from the "Management Center for IPS Sensors" page as shown in Figure 10.29 and follow these steps:

1. Select **Admin | Database**.
2. The "Database Rules" page appears. Select the radio button corresponding to the rule to view and click **View**.
3. The "View Database Rule" page appears. In the text box is detailed information about the rule. To return to the "Database Rules" page, click **OK**.

Deleting a Database Rule

To delete a database rule, start from the "Management Center for IPS Sensors" page as shown in Figure 10.29 and follow these steps:

1. Select **Admin | Database**.
2. The "Database Rules" page appears. Select the radio button corresponding to the rule you want to delete and click **Delete**. The database rule is deleted from the IPS Management Center.

Updating Sensor Software and Signatures

Cisco Systems is constantly providing new sensor software versions and signature release levels. These new versions and release levels are provided in files known as Service Pack update files and Signature update files.

The procedures to update the sensor software and the signatures are complex. To be informed of the latest update files by email, you can subscribe to the Cisco IPS Active Update Notification.

Defining the E-mail Server Settings

You can specify the email server that the Cisco IPS Management Center uses for event notification. To specify the server, follow these steps:

1. Start from the "Management Center for IPS Sensors" page as shown in Figure 10.29 and select **Admin | System Configuration**. Select **Email Server** in the Table of Contents.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:70

2. The "E-mail Server" page appears. Enter the email server name in the **Server Name** box. Click **Apply**. The email server specified will be used for event notification.

Snort

You have to love a program with the name Snort. This program is open source (it can be found at www.Snort.org), and there are mailing lists you can use to get support for it. Snort can be used as is, or you can use it with MySQL or add one of the many open source GUI interfaces to it. Whatever you choose, we show you how to use Snort effectively to defend your DMZ inexpensively. Formerly, when you installed Snort, you had to trudge through the installation, find the various dependencies, find the packages needed, install those packages, and a much later time, you would have the bare-bones IDS sensor running. We cover Snort's rules for signature matching and how Snort can use a few different formats for its log files.

Cisco has the name, but Snort has the price—it's *free*. When you are trying to protect your DMZ and conserve cash, free means a lot. But free in this case does not equate to a poor-quality product. Several companies have enhanced Snort with custom front ends and enhanced reporting functions.

One of the best ways to get a very clean packaged Snort with MySQL and a nice GUI interface is to download Eagle X from Engage Security at www.engagesecurity.com. The application is reasonably priced – free! It is quite easy to setup, since it installs and configures Apache, PHP, and MySQL for you. Figure A.53 shows the opening screen you see as you log into Eagle X. We assume you already know how to install and use Snort, but in this section of the chapter, we show you how to implement a Snort solution in your DMZ for intrusion detection.

NOTE

For more information on Snort, you can read *Snort 2.1 Intrusion Detection, Second Edition* (Syngress Publishing, Inc.; ISBN 1-931836-04-3) as well as *Snort Intrusion Detection and Prevention Toolkit* (Syngress Publishing, Inc.; ISBN 1-597490-99-7).

Eagle X gives a wealth of information up front without having to dig around for it. You get the signature match, sensor, timestamp, IP address information, and the event payload. Eagle X even has an online rule update process.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:71

Figure A.53 Eagle X Startup Screen

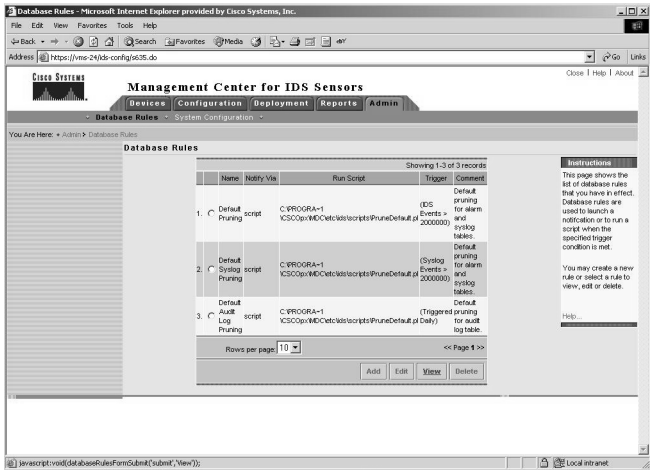

As with the Cisco IDS sensor we just read about, Snort IDS sensors follow the same deployment rules. If Snort cannot see the traffic, it cannot protect your network. Snort can be placed in the DMZ in the same manner as the Cisco sensor—one leg in front of the firewall or in the DMZ and the second leg in the trusted network. Unlike the Cisco sensor, Snort can function with only a single NIC installed in the sensor. The management and the sensing are done through the same NIC. However, this is not recommended, because you should not place an interface on the network being watched with a live IP address. Doing so would provide the attacker with a tip-off that there was something on the wire, and we want to be stealthy—not to mention that if the attacker could compromise your IDS, that is one less alarm he needs to worry about while he does his dirty deeds in your DMZ.

When configuring Snort, you need to be aware that some Snort distributions or scripts run Snort as root. This is not good security practice. Dan Bernstein has a tool package called *daemontools* that can be found at http://cr.yp.to/daemontools.html; it can help prevent this situation if you're running Snort on Linux-based hosts.

Once the assimilation with daemontools has been completed, Snort will not run as root in the daemon mode (*-D*). Keep in mind that several versions of Snort are available, so always try to get the newest one, which, at this time is 2.6. This version addresses several items; one of the most important provides a resolution for a memory leak issue. Having the latest code is a much preferred security posture on something as critical as the IDS sensor sitting in your DMZ. Snort will run and sense with only a single NIC installed, but again,

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:72

from a security perspective in the DMZ, this is not acceptable. Make sure to use at least *two* NICs for your Snort sensor. When configuring your Snort host, install SSH in order to provide a secure link from the sensor to a remote management or logging server. Make sure that all services such as FTP and X-Windows are disabled. We keep repeating: Leave no freebies for the would-be hacker of your DMZ.

Deploying Snort is no different than deploying any other NIDS, with one exception: It is one of the cheapest, so you can have a large number of sensors deployed around the network for a minimal amount of cash, covering places like the internal networks, the DMZ, and outside the DMZ. Combined with deployment of an HIDS such as Tripwire, Snort it is a very potent network defense against the wannabe hacker of your DMZ and network.

There are a variety of ways to manage Snort IDS sensors. An easy way to configure Snort is to use the Webadmin Snort plug-in found at www.snort.org/dl/contrib/front_ends/webmin_plugin/Figure A.54 shows a screen of the Snort IDS plug-in for managing Snort. From this plug-in we can set virtually all settings of Snort as well as the alerts and logging. We can also manage our rules for Snort from this tool.

Figure A.54 The Webadmin Snort Tool

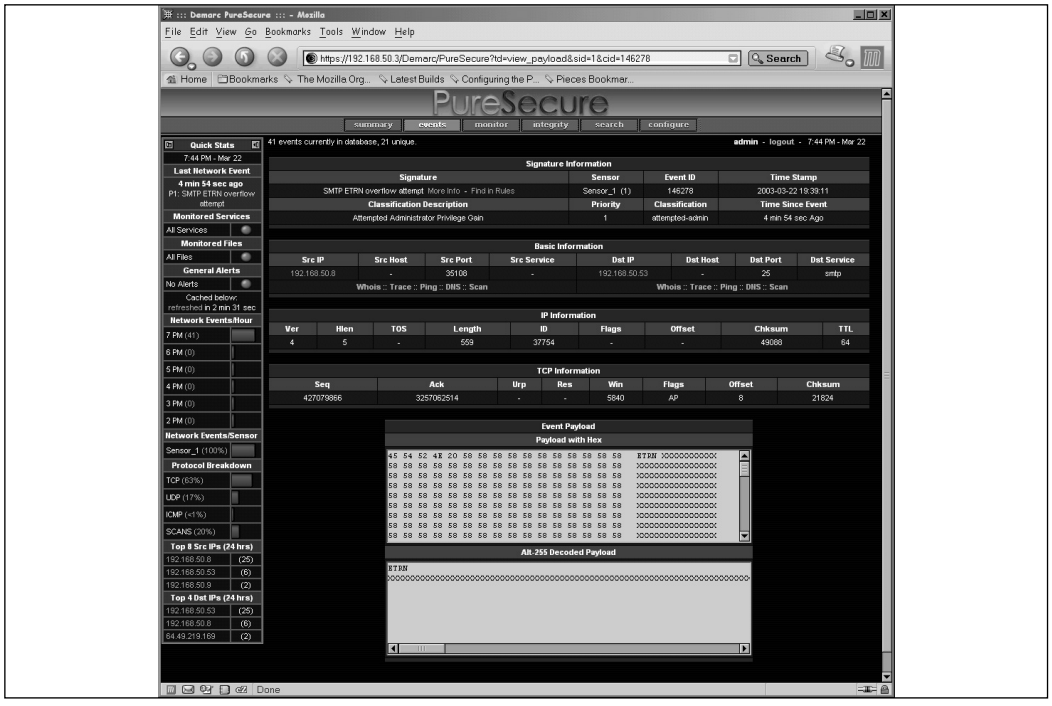

Another common database front-end for Snort is BASE (Basic Analysis and Security Engine). Smaller networks might enjoy the simplicity of "grepping" through their intrusion logs, but medium-sized and large enterprises need to rely on the structure of a well-main-

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:73

tained database. BASE is one of the best known open source analysis tools and in the next section we discuss installing, using, and maintaining it.

BASE

The Basic Analysis and Security Engine (BASE) is the successor of the Analysis Console for Intrusion Databases (ACID). Development around ACID has been stagnant for a long time and BASE fixes quite a few shortcoming of the tool. BASE, just as ACID, is a PHP-based analysis engine to manage a database of security events. These events can be from IDSs (such as Snort), but also from firewalls or network monitoring tools and even pcap files. The database schema that is used by BASE is based on the Snort database schema with some additional tables.

logsnorter, available at http://www.snort.org/dl/contrib/other_logs/logsnorter-0.2.tar.gz, is a script designed to import logs with alerts from these engines into Snort databases, so this data becomes available to BASE, too.

At this time, BASE provides the following features:

- An interface for database searching and query building. Searches can be performed by network-specific parameters such as attacker's IP address, by meta parameters such as time or date of an event, or by triggered rule.
- A packet browser that can decode and display Layer 3 and Layer 4 information from logged packets.
- Data management capabilities, including grouping of alerts (so that it is possible to group all events related to an intrusion incident), alert deletion, or archiving and exporting to e-mail messages.
- Generation of various graphical charts and statistics based on specified parameters.

The rest of this section describes the installation of BASE and its prerequisites, Snort configuration, and the ways in which BASE can be used for intrusion detection and analysis. You can download the latest release of BASE (1.2.6 as of July 2006) from http://sourceforge.net/projects/secureideas or install it from the accompanying CD-ROM.

Installing BASE

The structure of BASE is multi-tiered and scalable (see Figure A.55). You can use it on just one computer, or you can have architecture of up to three tiers.

NOTE

BASE is included on the accompanying CD-ROM.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:74

Figure A.55 Multi-tiered Architecture of an IDS and BASE Console

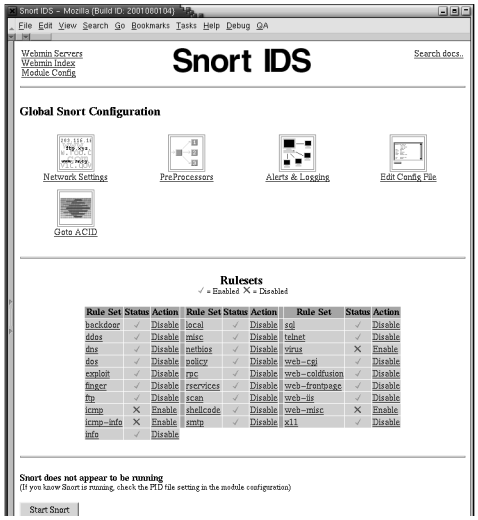
Snort IDS - Mozilla {Build ID: 2001080104}
File Edit View Search Go Bookmarks Tasks Help Debug QA
Webmin Servers
Webmin Index
Module Config
Snort IDS
Search docs..
Global Snort Configuration
Network Settings
PreProcessors
Alerts & Logging
Edit Config File
Goto ACID
Rulesets
√ = Enabled × = Disabled
Rule Set Status Action Rule Set Status Action Rule Set Status Action
backdoor √ Disable local √ Disable sql √ Disable
ddos √ Disable misc √ Disable telnet √ Disable
dns √ Disable netbios √ Disable virus × Enable
dos √ Disable policy √ Disable web-cgi √ Disable
exploit √ Disable rpc √ Disable web-coldfusion √ Disable
finger √ Disable rservices √ Disable web-frontpage √ Disable
ftp √ Disable scan √ Disable web-iis √ Disable
icmp × Enable shellcode √ Disable web-misc × Enable
icmp-info × Enable smtp √ Disable x11 √ Disable
info √ Disable
Snort does not appear to be running
(If you know Snort is running, check the PID file setting in the module configuration)
Start Snort

As you can see, BASE works with alerts stored in a database by sensors. A set of PHP scripts is used for creating queries and browsing the results. Currently, BASE officially supports PostgreSQL, MySQL and Microsoft SQL Server 2000, but it is possible to modify it to work with other SQL-based DBMSs supported by PHP. You can use any Web server as long as it supports PHP (although you might run into difficulties with the optional graphing functionality of BASE because the libraries it uses are mainly designed for Linux and Apache). As such, documentation provided with the application suggests MySQL as the database of choice.

NOTE

The decision of which OS to use is entirely up to you. Use the OS that you are most comfortable with, just *don't forget to harden it*. There are few things more embarrassing than finding out that one of your security systems has been compromised. Take time to make sure your BASE database and Web servers aren't going to be compromised.

Prerequisites for Installing BASE

Let's assume that a Web server and a database are installed on the same host. Your Snort sensor is probably located on another machine, although it is not important to us—BASE does not work directly with the sensor, only with the data reported into the database. If you would like to separate the Web server (front end) from the database (back end), almost

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:75

nothing changes in the BASE configuration—only some IP addresses in configuration files. It is even possible to have many Web servers working with one database. Moreover, of course, the number of Web clients is not limited even for one Web server.

Operating System on BASE Host

In this section, we mainly use Linux. The operating system used is not overly crucial; all the BASE components can be installed (with minimal modifications) on any UNIX operating systems or even Microsoft Windows (although the latter requires more tweaking). If you plan to use the BASE host only as a server, you can install a minimal set of packages—the only crucial parts are networking and software development tools. Actual selection of packages depends on your choice—it is easy to add any missing dependencies when they're needed.

We will set an IP address of 10.1.1.30 for our BASE server.

Tools & Traps…

When Size Matters

Running Snort on a busy network can produce a significant number of alerts. With a standard set of rules, it can generate tens of megabytes of data per day on a network with just a couple of busy Web sites. Additionally, nothing stops you from writing configuration files for logging interesting data to store as a reference for future investigations. This data can quickly fill a hard drive.

If you have only one partition that holds the entire file system, filling it up might cause the machine to stop functioning. It is considered good practice to separate the log and database partitions from the / (root) and /boot partitions.

The Web Server

We will use the Apache 2.0.x Web server on Linux because it is a native environment for BASE. You can either download it from www.apache.org and compile manually or use a package that comes with your Linux distribution. For example, to install Apache on a Debian system, use the following commands:

```
# apt-get install apache2
# /etc/init.d/apache2 start
```

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:76

These commands install the package and automatically add it to the daemons started by default.

PHP

BASE scripts are written in PHP language, so naturally we need to add PHP4 support to our Web server. There are many different ways to set it up. For example, it can be set up as an Apache module or run as an external CGI application. The important features for us are:

- Database support - MySQL, PostgreSQL, or MS SQL. We use MySQL throughout this section.
- GD support - This is a graphing library used for producing graphs.
- Socket support - This is used only for performing native whois queries.

You can either build PHP from source or use pre-compiled packages for your system. When building from source, you need to use at least the following options in PHP configuration. For MySQL support:

```
./configure [your config options] --with-mysql --with-gd --enable-sockets
```

For PostgreSQL support:

```
./configure [your config options] --with-pgsql --with-gd --enable-sockets
```

Using the option *—with-apache* makes PHP work as an Apache Web server module and in doing so, significantly speeds up script execution. If you do not want to deal with compiling the source, it is possible to use Linux packages that are already included in the distribution. Their names vary from distribution to distribution. In Debian you can install them as follows:

```
# apt-get install php4 php4-mysql php4-gd
```

After installation, it is recommended that you modify the php.ini configuration file:

1. Disable display of inline PHP error messages in generated HTML files by setting display errors=off in production environment or at least set error reporting = E_ALL & ~E_NOTICE, which will limit the number of reported error messages.
2. Configure SMTP on the server. On Windows you need to set the SMTP variable to the path of your SMTP server executable module. On UNIX, set the sendmail_path to the path of the sendmail executable (for example, sendmail_path=/usr/sbin/sendmail).
3. On Windows platforms you also need to set the session.save_path variable to a temporary directory writable by the Web server (for example, c:\temp). Windows-related configuration and installation issues are documented at www.php.net/manual/en/install.windows.php.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:77

Support Libraries

The following libraries need to be installed. Not all of them are critical for BASE functionality. The only important one is ADODB; others can be omitted if you are ready to sacrifice graphing features of BASE.

We already mentioned the GD library. This library for raw image manipulation supports GIF/JPEG/PNG formats. It is available at www.boutell.com/gd. The minimum version that can be used with BASE is 1.8. GD depends on some other libraries (usually installed as a part of system setup, but just in case we list them here):

- libpng, available at www.libpng.org/pub/png
- libjpeg-6b, available at www.ijg.org
- zlib, available at www.gzip.org/zlib

Manually installing ADODB, which is available at http://adodb.sourceforge.net/, is done as follows:

```
$ cp adodb491.tgz /var/www/html
$ cd /var/www/html
$ tar -xvzf adodb491.tgz
$ mv adodb491 adodb
```

If you are using Debian, the installation is done with simply executing this command:

```
$ sudo apt-get install libphp-adodb
```

MySQL or PostgreSQL

The database to gather the Snort events is probably already installed; you simply need to follow general recommendations for setting up database logging with Snort. If it is not installed, you can use the packages from your Linux distribution or download them from www.mysql.com. The setup of database logging is described in Section 2.3.6 of the Snort User's Manual, found at www.snort.org/docs/snort_htmanuals/htmanual_260/node13.html. We assume that Snort is set up to log in to MySQL database called snort_db, which is located on the same host as the Web server. The MySQL user used for logging is snort, and the password is password. You can use other values; just make sure that you set up proper permissions for database users. The Snort configuration file snort.conf must have the following line to log in to our database:

```
output database: log, mysql, user=snort password=password dbname=snort_db
host=10.1.1.30
```

Database tables need to be set up properly. A script create_mysql is included in the Snort distribution (in /contrib subdirectory, also there is one for PostgreSQL setup); when run, this script creates all necessary tables. The script can be run as follows:

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:78


```
# mysqladmin –u root –p create snort_db
# mysql –u root -p<password> snort_db < create_mysql
```

Next, create two users (snort for allowing the Snort sensor to log in to database and base for the BASE console to manipulate the data in the same database) and set passwords for them. You can (and should) omit the DELETE privilege here so the corresponding user will not be able to delete records from the database. For example, you can create a copy of the BASE console that will work under the user account that can browse events but not delete them.

```
mysql> grant INSERT, SELECT on snort_db.* to snort;
mysql> grant INSERT, SELECT on snort_db.* to snort@%;
mysql> grant CREATE, INSERT, SELECT, DELETE, UPDATE on snort.* to base;
mysql> grant CREATE, INSERT, SELECT, DELETE, UPDATE on snort.* to base@%;
```

Finally, set passwords for these users:

```
mysql> connect mysql
mysql> set password for 'snort'@'localhost' = password('password');
mysql> set password for 'snort'@'%' = password('password');
mysql> set password for 'base'@'localhost' = password('basepassword');
mysql> set password for 'base'@'%' = password('basepassword');
mysql> flush privileges;
mysql> exit
```

Note that without the *flush privileges* command, changes in password and privilege settings will NOT become effective.

Activating BASE

Actually installing BASE is simple too. You need to put the set of scripts in a location under the Web server root directory, for example:

```
$ cp base-1.2.6.tar.gz /var/www/html
$ cd /var/www/html
$ tar -xvzf base-1.2.6.tar.gz
```

Now that we are finished installing packages, let's proceed to BASE configuration.

Configuring BASE

First we need to set up some parameters for BASE to work with the database. The main configuration file for BASE is a base_conf.php.dist file located in the BASE directory on the Web server. This file needs to be copied to base_config.php:

```
$ cp base_conf.php.dist base_config.php
```

Table A.1 lists the most important parameters that can be defined in the configuration file.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:79

Table A.5 BASE Database Configuration Parameter

$DBlib_path	Full path to the ADODB installation (Note: Do not include a trailing \ character in any of the path variables) If you are using the Debian way of installing ADODB, it will be installed in /usr/share/php/adodb and not in your Web documents root. You will need to update your BASE configuration to reflect this path.
$Dbtype	Type of the database used (mysql, postgres, mssql)
$alert_dbname	Alert database name
$alert_host	Alert database server
$alert_port	Port on which MySQL, PostgreSQL or MSSQL server is listening (no need to change it if the default port is used)
$alert_user	Username for the alert database
$alert_password	Password for the username

In our case, they are configured as follows:

```
$DBlib_path = "/usr/share/php/adodb"
$DBtype = "mysql"
$alert_dbname = "snort_db"
$alert_host ="10.1.1.30"
$alert_user ="base"
$alert_password ="basepassword"
```

Another set of database parameters can be used for archiving alerts (moving them from the active database to a backup one):

- $archive_exists set to 1 to enable the feature.
- $archive_dbname Archive/backup database name
- $archive_host Archive database server
- $archive_port Port number for archive database server
- $archive_user Username for archive database
- $archive_password Password for this username

It is always a good idea to protect access to the BASE pages with a Web server password. As an example, we will require a username *admin* and password *adminpassword* from a user trying to access the location */base* on a Web server via the Web browser:

```
# mkdir /usr/lib/apache/passwords
# htpasswd -c /usr/lib/apache/passwords/.htpasswd admin
(enter "adminpassword" at the prompt)
```

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:80

Then the following lines need to be added to the httpd.conf file—a configuration file for the httpd daemon. In Debian, this file is located in the /etc/apache directory.

```
<Directory "/var/www/html/base">
AuthType Basic
AuthName "BASE console"
AuthUserFile /usr/lib/apache/passwords/.htpasswd
Require user admin
AllowOverride None
</Directory>
```

After making these changes, you need to restart the httpd daemon:

```
/etc/init.d/apache2 restart
```

Now we are ready to connect to the console for the first time. Accessing the URL http://10.1.1.30/base first brings up a request for a password, and then a page indicating that BASE has not been configured yet, meaning that some database tables for BASE are missing. BASE adds some extra tables to the database. Clicking the link **Setup page** runs a script that updates the database with the required tables. After clicking the Create BASE AG button on the next page, we are ready to start using BASE.

Damage & Defense…

BASE Security

As you probably noticed, no security features are embedded in BASE itself; therefore, to ensure security of its setup, you need to do additional tweaking. Your requirements will determine which tools you will use.

For one, you might be interested in using SSL (HTTPS connections) or TLS instead of plaintext communications between the browser and the server. In Apache, this is achieved using the mod_ssl module (www.modssl.org).

As you have previously seen, access to the BASE console can be restricted using native Web server authentication mechanisms—passwords or certificates. As was also previously mentioned, it might be useful to create at least two separate copies of BASE and configure one of them with only read database permissions. To restrict permissions for a specific copy of BASE, simply revoke the *DELETE* privilege from the database user configured in this copy.

The most important security issue is that all database passwords are hard-coded in the PHP scripts in cleartext, so extreme caution needs to be applied to the host configuration. Any exposure of source code for PHP scripts will expose the password to an attacker.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:81

Using BASE

Using BASE is rather simple. Its screens are self-explanatory most of the time. Let's look at the main screen (see Figure A.56).

Figure A.56 The BASE Main Screen

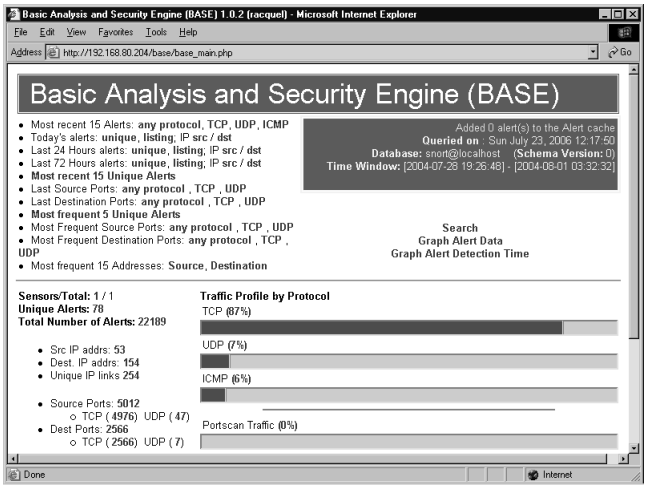
Basic Analysis and Security Engine (BASE) 1.0.2 (racquel) - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://192.168.80.204/base/base_main.php
Go
Basic Analysis and Security Engine (BASE)
Most recent 15 Alerts: any protocol, TCP, UDP, ICMP
Today's alerts: unique, listing, IP src / dst
Last 24 Hours alerts: unique, listing, IP src / dst
Last 72 Hours alerts: unique, listing, IP src / dst
Most recent 15 Unique Alerts
Last Source Ports: any protocol , TCP , UDP
Last Destination Ports: any protocol , TCP , UDP
Most frequent 5 Unique Alerts
Most Frequent Source Ports: any protocol , TCP , UDP
Most Frequent Destination Ports: any protocol , TCP , UDP
Most frequent 15 Addresses: Source, Destination
Added 0 alert(s) to the Alert cache
Queried on : Sun July 23, 2006 12:17:50
Database: snort@localhost (Schema Version: 0)
Time Window: [2004-07-28 19:26:48] - [2004-08-01 03:32:32]
Search
Graph Alert Data
Graph Alert Detection Time
Sensors/Total: 1 / 1
Unique Alerts: 78
Total Number of Alerts: 22189
Src IP addrs: 53
Dest. IP addrs: 154
Unique IP links 254
Source Ports: 5012
TCP (4976) UDP (47)
Dest Ports: 2566
TCP (2566) UDP (7)
Traffic Profile by Protocol
TCP (87%)
UDP (7%)
ICMP (6%)
Portscan Traffic (0%)
Done
Internet

This screen shows the general statistics for BASE; namely, the number of alerts divided by protocol, the counts of source and destination ports for triggered rules, and so forth. Clicking a link provides additional details about the particular category. Figure A.57 provides an example listing of all the unique alerts (alerts grouped by the triggered rule).

Figure A.57 Unique Alerts

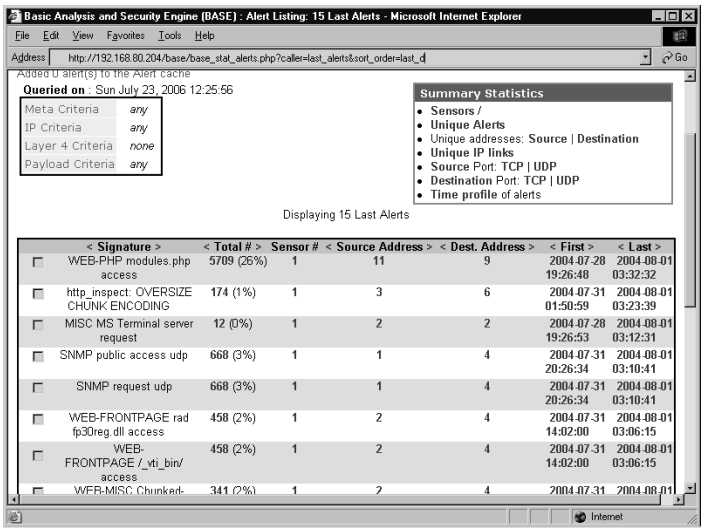
Basic Analysis and Security Engine (BASE) : Alert Listing: 15 Last Alerts - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://192.168.80.204/base/base_stat_alerts.php?caller=last_alerts&sort_order=last_d
Go
Added 0 alert(s) to the Alert cache
Queried on : Sun July 23, 2006 12:25:56
Meta Criteria any
IP Criteria any
Layer 4 Criteria none
Payload Criteria any
Summary Statistics
Sensors /
Unique Alerts
Unique addresses: Source | Destination
Unique IP links
Source Port: TCP | UDP
Destination Port: TCP | UDP
Time profile of alerts
Displaying 15 Last Alerts
< Signature > < Total # > Sensor # < Source Address > < Dest. Address > < First > < Last >
WEB-PHP modules.php access 5709 (26%) 1 11 9 2004-07-28 19:26:48 2004-08-01 03:32:32
http_inspect: OVERSIZE CHUNK ENCODING 174 (1%) 1 3 6 2004-07-31 01:50:59 2004-08-01 03:23:39
MISC MS Terminal server request 12 (0%) 1 2 2 2004-07-28 19:26:53 2004-08-01 03:12:31
SNMP public access udp 668 (3%) 1 1 4 2004-07-31 20:26:34 2004-08-01 03:10:41
SNMP request udp 668 (3%) 1 1 4 2004-07-31 20:26:34 2004-08-01 03:10:41
WEB-FRONTPAGE rad fp30reg.dll access 458 (2%) 1 2 4 2004-07-31 14:02:00 2004-08-01 03:06:15
WEB-FRONTPAGE /_vti_bin/ access 458 (2%) 1 2 4 2004-07-31 14:02:00 2004-08-01 03:06:15
WEB-MISC Chunked- 341 (2%) 1 2 4 2004-07-31 2004-08-01
Internet

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:82

Each line (alert) has several clickable fields; the most interesting of these are probably the classification field and the references to various attack databases links (for example, Arachnids or CVE). This data is taken from rules when Snort logs an alert to the database. If you click the **cve** link in the line that has such a link in the *Signature* field, you will be taken to the description of this attack in CVE (a database of vulnerabilities). The Snort link leads to the similar description on the www.snort.org site. Classification helps group attacks by their type, which is also set up in the Snort rules file.

Each individual logged packet can be displayed in a decoded format, showing various flags, options, and packet contents (see Figure A.58).

Figure A.58 Displaying a Single Alert

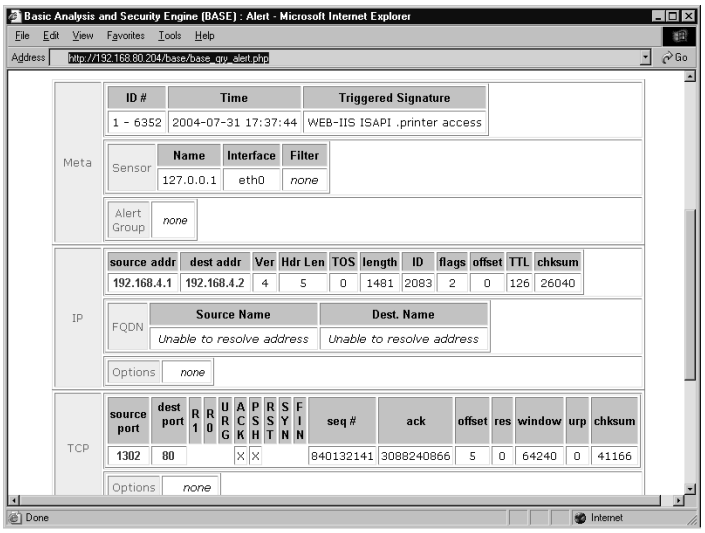
Basic Analysis and Security Engine (BASE) : Alert - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://192.168.80.204/base/base_qry_alert.php
Go
Meta
ID # | Time | Triggered Signature
1 - 6352 | 2004-07-31 17:37:44 | WEB-IIS ISAPI .printer access
Sensor
Name | Interface | Filter
127.0.0.1 | eth0 | none
Alert Group none
IP
source addr | dest addr | Ver | Hdr Len | TOS | length | ID | flags | offset | TTL | chksum
192.168.4.1 | 192.168.4.2 | 4 | 5 | 0 | 1481 | 2083 | 2 | 0 | 126 | 26040
FQDN
Source Name | Dest. Name
Unable to resolve address | Unable to resolve address
Options none
TCP
source port | dest port | R 1 | R 0 | URG | ACK | PSH | RST | SYN | FIN | seq # | ack | offset | res | window | urp | chksum
1302 | 80 | | | | X | X | | | | 840132141 | 3088240866 | 5 | 0 | 64240 | 0 | 41166
Options none
Done
Internet

The unique alert display can be used for checking any "noisy" signatures and tuning them. You can sort the listing in ascending or descending order of number of alerts and then select the ones that are triggered more often. Sorting is done by clicking a corresponding arrow (**>** or **<**) in the header of the relevant column (refer back to Figure A.35).

Querying the Database

One of the most important features of BASE is its searching tools. It is possible to create database queries with many parameters—from signature type to packet payload contents (provided that this information has been logged in the database). The main search screen is shown in Figure A.59.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:83

Figure A.59 Search Parameters

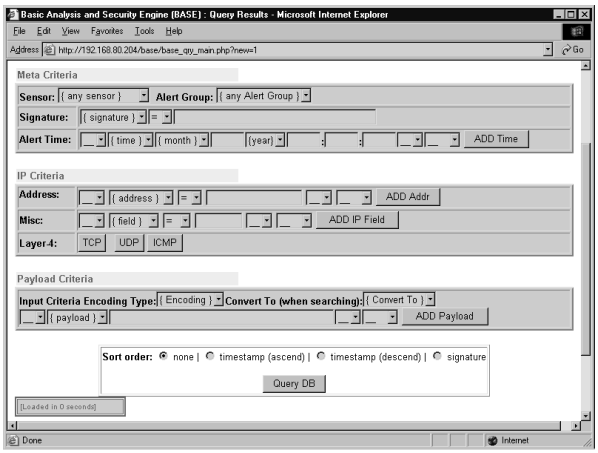
Basic Analysis and Security Engine (BASE) : Query Results - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://192.168.80.204/base/base_qry_main.php?new=1
Meta Criteria
Sensor: { any sensor }
Alert Group: { any Alert Group }
Signature: { signature } =
Alert Time: { time } { month } (year) ADD Time
IP Criteria
Address: { address } = ADD Addr
Misc: { field } = ADD IP Field
Layer 4: TCP UDP ICMP
Payload Criteria
Input Criteria Encoding Type: { Encoding } Convert To (when searching): { Convert To }
{ payload } ADD Payload
Sort order: none | timestamp (ascend) | timestamp (descend) | signature
Query DB
[Loaded in 0 seconds]
Done
Internet

As you can see, in the Meta Criteria section, you can specify different Snort sensors (in a case where you have many sensors storing data in the same database), search in a specific alert group only (more about alert groups in the next section), and match signatures (exactly or by a substring in their names), classification, and time periods. It is also possible to search only for packets with specific Layer 3 and Layer 4 information, plus perform a context search inside captured packets' payload. For example, let's find all alerts triggered by signatures related to Nmap scanner. This can be achieved by specifying the *signature* field in *meta criteria* as *roughly = NMAP* and clicking the **Query DB** button. The result of this query is shown in Figure A.60.

Figure A.60 All NMAP-Related Alerts from the Database

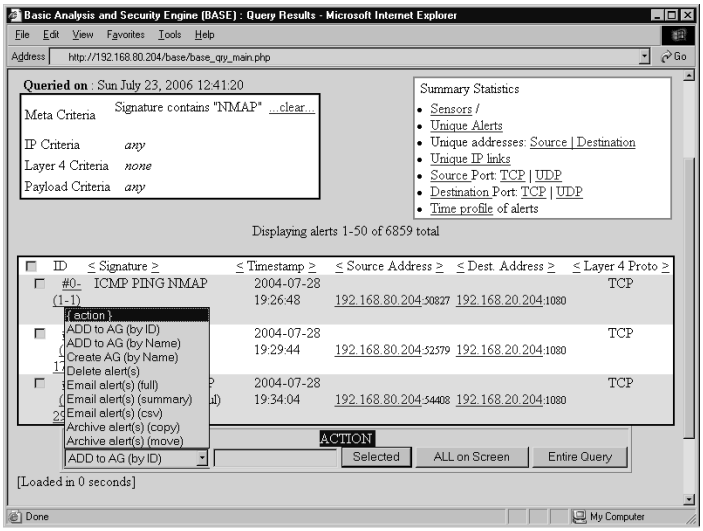
Basic Analysis and Security Engine (BASE) : Query Results - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://192.168.80.204/base/base_qry_main.php
Queried on : Sun July 23, 2006 12:41:20
Meta Criteria Signature contains "NMAP" ...clear...
IP Criteria any
Layer 4 Criteria none
Payload Criteria any
Summary Statistics
Sensors / Unique Alerts
Unique addresses: Source | Destination
Unique IP links
Source Port: TCP | UDP
Destination Port: TCP | UDP
Time profile of alerts
Displaying alerts 1-50 of 6859 total
ID < Signature > < Timestamp > < Source Address > < Dest. Address > < Layer 4 Proto >
#0-(1-1) ICMP PING NMAP 2004-07-28 19:26:48 192.168.80.204:50827 192.168.20.204:1080 TCP
2004-07-28 19:29:44 192.168.80.204:52579 192.168.20.204:1080 TCP
2004-07-28 19:34:04 192.168.80.204:54408 192.168.20.204:1080 TCP
{ action }
ADD to AG (by ID)
ADD to AG (by Name)
Create AG (by Name)
Delete alert(s)
Email alert(s) (full)
Email alert(s) (summary)
Email alert(s) (csv)
Archive alert(s) (copy)
Archive alert(s) (move)
ACTION
ADD to AG (by ID)
Selected
ALL on Screen
Entire Query
[Loaded in 0 seconds]
Done
My Computer

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:84

In the bottom-left corner is an *action* field, which specifies possible actions that you can perform with the results of the query. The displayed alerts can be added to an alert group, deleted from the database, e-mailed in various formats, or archived to another database. The three buttons on the right specify which alerts are used when the selected action is performed. If you click the Selected button, only specifically selected alerts from all displayed will be used (the leftmost column of the table contains check boxes for row selection). If you click the ALL on Screen button, all displayed alerts are used, and clicking the Entire Query button uses the entire set of results. The difference between ALL on Screen and Entire Query is that when many results are returned, they are displayed in sets of 50.

The Email alerts action takes as a parameter an address where the results should be sent. This address is entered in a provided field. The Add to AG action also takes a parameter—an *alert group* name or number. Other actions do not need parameters.

Actually, almost all the buttons on the front page of the BASE console are simply shortcuts for various queries that could be constructed via the main search interface.

Alert Groups

Alert groups are entities used to logically group various alerts and attach annotations to sets of events (incidents). An alert group has a number, a text name, and an optional annotation or commentary. For example, if you are researching a particular intrusion incident, you might be interested in putting all the related alerts into one group so that you will be able to reference it in running queries, e-mailing results, and so forth. To do the grouping, you need to create the group first. When you click the link **Alert Group (AG) Maintenance** at the bottom of the BASE main screen, you are presented with the window shown in Figure A.61.

Figure A.61 Listing of Alert Groups

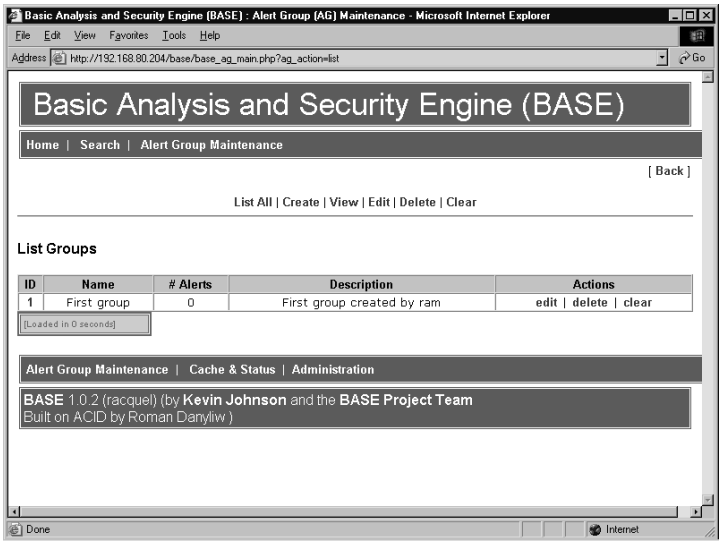
Basic Analysis and Security Engine (BASE) : Alert Group (AG) Maintenance - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://192.168.80.204/base/base_ag_main.php?ag_action=list
Go
Basic Analysis and Security Engine (BASE)
Home | Search | Alert Group Maintenance
[Back]
List All | Create | View | Edit | Delete | Clear
List Groups
ID
Name
Alerts
Description
Actions
1
First group
0
First group created by ram
edit | delete | clear
[Loaded in 0 seconds]
Alert Group Maintenance | Cache & Status | Administration
BASE 1.0.2 (racquel) (by Kevin Johnson and the BASE Project Team
Built on ACID by Roman Danyliw)
Done
Internet

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:85

In our example, we are using the ID of *1* and the name *first group*. To create another group, click the **Create** link at the top of this page. You will be asked to enter the name for the new group and an optional description. For our example, we used *grinder incident* as the name of the new group. The group ID is generated automatically. When this information is saved, the list of groups appears similar to the window shown in Figure A.62.

Figure A.62 Creating a New Group

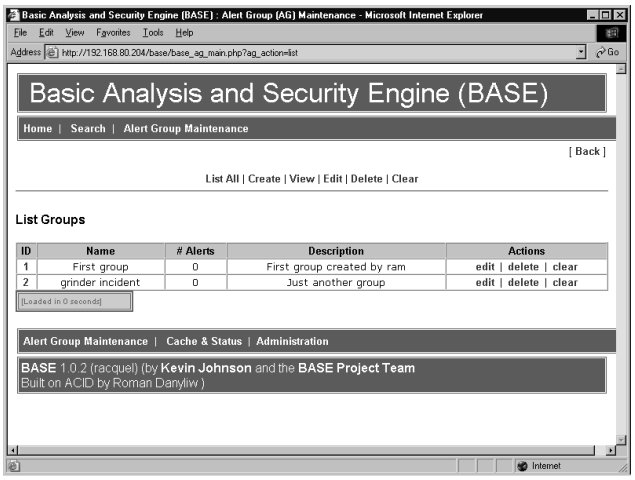
Basic Analysis and Security Engine (BASE) : Alert Group (AG) Maintenance - Microsoft Internet Explorer
Basic Analysis and Security Engine (BASE)
Home | Search | Alert Group Maintenance
[Back]
List All | Create | View | Edit | Delete | Clear
List Groups
ID Name # Alerts Description Actions
1 First group 0 First group created by ram edit | delete | clear
2 grinder incident 0 Just another group edit | delete | clear
Alert Group Maintenance | Cache & Status | Administration
BASE 1.0.2 (racquel) (by Kevin Johnson and the BASE Project Team
Built on ACID by Roman Danyliw)

Now we can run a query (for example, let's search for all NMAP-related alerts again) and add the results to Group 2. When presented with the query results, select an action **Add to AG (by ID)** and enter **2** as an ID. Alternatively, you can use **ADD to AG (by Name)** and enter the name given to our group. After you click **Entire Query**, all search results will be added to the specified group. Figure A.64 shows how the parameters are entered in the Query Results screen, and Figure A.63 displays the resulting listing of the groups.

Figure A.63 Adding Search Results to an Alert Group

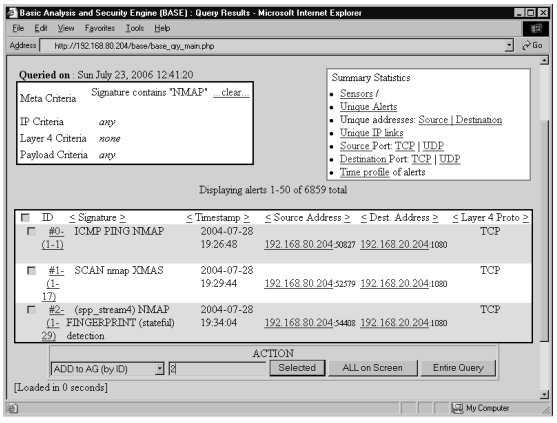
Basic Analysis and Security Engine (BASE) : Query Results - Microsoft Internet Explorer
Queried on : Sun July 23, 2006 12:41:20
Meta Criteria Signature contains "NMAP" ...clear...
IP Criteria any
Layer 4 Criteria none
Payload Criteria any
Summary Statistics
Displaying alerts 1-50 of 6859 total
ACTION
ADD to AG (by ID)
Selected
ALL on Screen
Entire Query
[Loaded in 0 seconds]

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:86

Figure A.64 Result of Alert Grouping

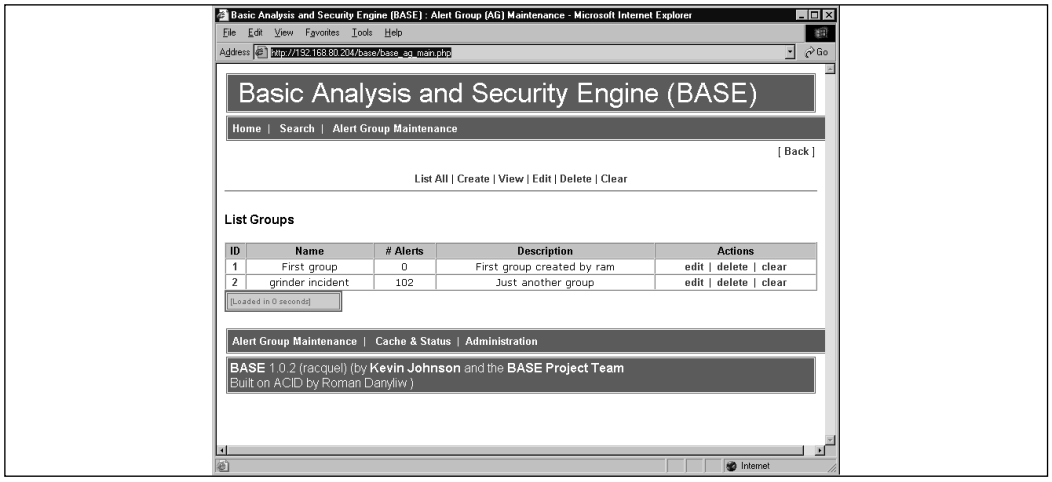
Basic Analysis and Security Engine (BASE) : Alert Group (AG) Maintenance - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://192.168.80.204/base/base_ag_main.php
Go
Basic Analysis and Security Engine (BASE)
Home | Search | Alert Group Maintenance
[Back]
List All | Create | View | Edit | Delete | Clear
List Groups
ID Name # Alerts Description Actions
1 First group 0 First group created by ram edit | delete | clear
2 grinder incident 102 Just another group edit | delete | clear
[Loaded in 0 seconds]
Alert Group Maintenance | Cache & Status | Administration
BASE 1.0.2 (racquel) (by Kevin Johnson and the BASE Project Team
Built on ACID by Roman Danyliw)
Internet

Each group can be modified:

- The Edit link presents you with the screen for modifying the group's name and description.
- The Delete link deletes the group. It does not delete the alerts, only the group as a logical entity.
- The Clear link clears a group's contents by un-grouping all alerts from it; it does not delete the alerts from the database.

NOTE

An alert can be part of multiple groups simultaneously.

Graphical Features of BASE

BASE has a tool that can produce a graphical summary of alerts based on date periods, alert group membership, source and destination ports, and IP addresses. An interface for the graph generation is shown in Figure A.65.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:87

Figure A.65 Alert Graphing

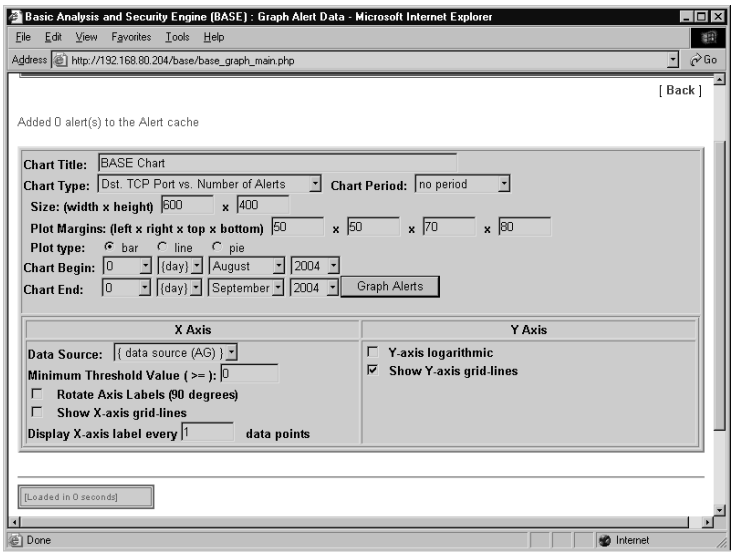
Basic Analysis and Security Engine (BASE) : Graph Alert Data - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Address http://192.168.80.204/base/base_graph_main.php
Go
[Back]
Added 0 alert(s) to the Alert cache
Chart Title: BASE Chart
Chart Type: Dst. TCP Port vs. Number of Alerts
Chart Period: no period
Size: (width x height) 600 x 400
Plot Margins: (left x right x top x bottom) 50 x 50 x 70 x 80
Plot type: bar line pie
Chart Begin: 0 (day) August 2004
Chart End: 0 (day) September 2004
Graph Alerts
X Axis
Data Source: (data source (AG))
Minimum Threshold Value (>=): 0
Rotate Axis Labels (90 degrees)
Show X-axis grid-lines
Display X-axis label every 1 data points
Y Axis
Y-axis logarithmic
Show Y-axis grid-lines
[Loaded in 0 seconds]
Done
Internet

Many of the features within the graph parameters are relatively self-explanatory:

- The Chart Type parameter allows for the selection of a specific type of graph to be generated.
- The Data Source parameter allows limiting alerts by date, specified by the Chart Begin and Chart End parameters, and by alert group. If you select an alert group in this drop-down box, only alerts from this group will be used as a source dataset.

Another interesting feature is the *Chart Period* parameter. If nothing is selected here, the X axis will list either all dates or all ports/IPs, depending on the chart type. If you select a period such as a week or a day, all alerts are grouped by day of the week or hour of the day. This allows creation of statistics such as daily distribution of alerts depending on a day of the week or time of day. Try it, and you will see that most attacks usually happen during the night and/or on weekends (at least the script kiddies' attacks, which amount to the biggest part of intrusion traffic). Figure A.66 shows a sample BASE chart.

Managing Alert Databases

The database of alerts produced by Snort sensors grows with time. If a significant number of alerts are logged, the database will become quite large, resulting in slow searches. To keep the alert database to a manageable size, you can use a variety of methods.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:88

Figure A.66 A Sample BASE Chart

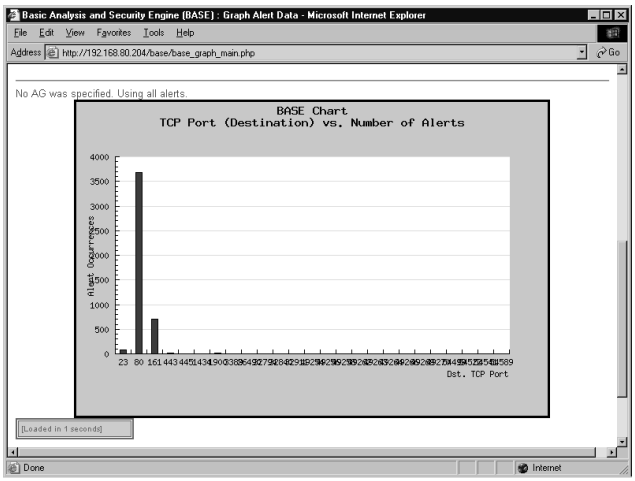

The simplest management technique is referred to as *trimming*. Simply put, trimming translates to deleting the uninteresting and older alerts triggered by false positives. If you want to delete an alert or a set of alerts, run a query that includes the alert as one of the results, choose the **Delete Alerts** action in the Results screen, and press the corresponding button:

- Click Selected if you want to delete only part of alerts displayed.
- Click All on Screen to delete all displayed alerts.
- Click Entire Query to delete all results of the current query.

Another management technique is *archiving*. Archiving is the process by which you move the undesired alerts to another database. To use this feature, you need to create a second database in exactly the same way that the main one was created. This is accomplished using the *create_mysql* or *create_postgresql* scripts (for information on how to use these scripts, review the section "Installing BASE Prerequisites" in this chapter). Let's assume that this database is called *snort_archive*. After that, you need to specify parameters of this database in *base_conf.php* file; for example:

```
$archive_dbname = "snort_archive"
$archive_host ="10.1.1.30"
$archive_user ="base"
$archive_password ="basepassword"
```

Now after running a query it is possible to select an action **Archive alerts (move)** or **Archive alerts (copy)**. After one of the buttons **Selected**, **ALL on Screen**, or **Entire Query** is pressed, corresponding alerts are moved (or copied) in the archive database. You

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:89

can set up a second copy of BASE in another Web server directory and specify this archive database as active for this copy. After that, you will be able to browse the archive as well.

To sum up, BASE is one of the most mature open source GUI tool for interactive Snort event analysis.

One of the joys of Snort is that you can write your own rules for it. Doing so is not really for the faint of heart, but it does have immense value. You can tweak an existing rule for your own purposes or come up with a completely custom rule based on your needs. As with the Cisco IDS, you can finely tune Snort in the DMZ to look for specific intrusions based on the hosts in the DMZ. You enable or disable rules, make rules with specific Boolean logic, IP addressing, ports, or specific strings. Snort.org has a comprehensive rule database online where you can search for a specific rule or find something close and then modify it to suit your needs. A second resource can be found at www.bleedingsnort.com. They have matched the rules with packet trace files and the research that accounted for the rule. It's a great way to see how and why a rule was coded in a certain manner. In Figure A.67 we see a sample rule for Snort. Since Snort does not understand multiple lines, the rules consist of a single line.

Figure A.67 Sample Snort Rule

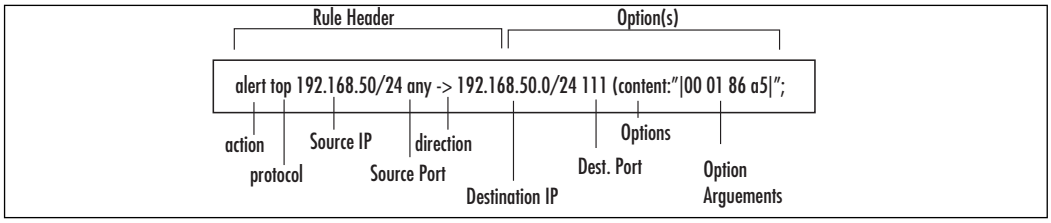

The text in the rule up to the first parenthesis is the *rule header:*

```
alert tcp 192.168.50.0/24 any -> 192.168.1.0/24 111 (……)
```

The first item is the *rule action*. In this example, it is *alert*. This is what Snort is supposed to do when Snort finds a match to this rule. There are five choices of action to choose from:

- **Alert** Generates an alert using the selected alert method and then logs the packet.
- **Log** Logs the packet.
- **Pass** Drops or ignores the packet.
- **Activate** Alerts and then activates another dynamic rule
- **Dynamic** Do nothing until enabled by an "activate" rule, then run as a log rule

The next item is the *protocol field:*

```
alert <protocol> 192.168.50.0/24 any -> 192.168.1.0/24 111 (……)
```

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:90

There are four protocols that Snort can watch for:

- tcp
- udp
- icmp
- ip

However, with future versions, these numbers will more than likely change.

Next in line is the *IP address field:*

```
alert tcp <IP address> 111 (……)
```

We can specify specific IP address, wildcard addresses using *any*, or a range of IP addresses using the CIDR block mask. This is where we specify a subnet mask by the slash (/) and then the number of bits to mask the address. An example is 192.168.50.0 /24, which translates to 255.255.255.0.

The final field before the parenthesis is the *port field*.

```
alert tcp 192.168.50.0/24 any -> 192.168.1.0/24 <port> (……)
```

We can specify a certain port, a range of ports, or any port.

- **1:1024** Tells the rule to use ports 1 to 1024.
- **:6000** Tells the rule to use any ports less than or equal to 6000.
- **400:** Tells the rule to use any ports greater than or equal to 400.

We can tell the rule to ignore something by using the negation (colloquially known as the *bang*) or ! symbol. This tells the rule to do something to ports except what comes after the ! symbol.

In the rule we can tell Snort which direction to watch by using the *directional operator* (->) in the rule, or we can say "any direction" using the <> operator.

Since this chapter is not about writing rules for Snort, we leave the options to you to study. The options can be quite simple or very complex, with offsets and Boolean logic. For more detailed information on writing Snort rules, you can find the information in Syngress's book on Snort that we mentioned earlier in this chapter or by going to the Snort Web site.

Many Snort enhancements are available, ranging from graphical interfaces, signature generators, and log file analyzers to plug-ins. Snort can be one of the most valuable tools to protect your DMZ and network from intruders and attacks. For the price, Snort is one of the best bargains for network security and IDS systems.

The Poor Man's IDS

Until the advent of Snort and other open-source IDS, you had to be prepared to spend a significant amount of money on an IDS. Or perhaps you needed an IDS in the DMZ for a

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:91

short amount of time, say to identify Kazaa or a suspected Web hack taking place or to look for a worm like CodeRed. What can you do with existing tools and not break the budget?

In many packet-sniffing tools, you can design and set up custom filters and then have the packet analyzer look for the match and trip an alarm. That way, you could have the analyzer running and checking for Kazaa, Back Oriface, various worm signatures, and so on. The signatures are pretty easy to come by at sites such as the Snort Web site rule database or the Snort signature mail list at https://lists.sourceforge.net/lists/listinfo/snort-sigs. The signatures give you the details you need to look for in the packets. In the example in Figure A.68, we can see the signature string for the Blahot worm:

Figure A.68 Signature of Blahot Worm for Snortalert

```
tcp $HOME_NET any -> $EXTERNAL_NET $HTTP_PORTS (msg: "EXPLOIT Blahot Worm
Infection Reporting in"; flow: to_server,established;
uricontent:"/scr2/command.php?IP="; nocase; uricontent:"Port1="; nocase;
classtype: trojan-activity; reference:url,www.vitalsecurity.org/2005/01/malware-
spam.html; reference:url,www.blahot.com; sid: 2001667; rev:7; )
```

The part of the signature we are interested in is *uricontent: **"/src2/command.php?IP="***. This string is what we want our custom packet filter to look for. Depending on what exactly you want to look for, you might need to calculate the offset in the packet in order to tell the analyzer just where to look for the data. Some packet analyzers such as NAI Sniffer and Etherpeek let you use Boolean logic when you build the filter. These can be a real help when you want to find, say, FTP traffic but the FTP traffic is not on the expected port. You can use the custom filter to find all packets that contain FTP strings but *not* look at packets using the expected port 21. This way you could look for FTP commands as strings but filter out the normal FTP traffic on your network.

It's relatively easy to make a custom filter in Etherpeek. You start by selecting either **Edit** and then **Insert**, as shown in Figure A.69, or right-clicking the filters and then choosing **Insert**.

In this box, we need to click **Simple** and change it to **Advanced**. This will bring up the Boolean logic editor that we want. In Figure A.71 we see that the title is Port 80 Attack Fingerprints. We use Boolean logic to look for four different strings that would signal a port attack on a Web server. To add another logic function, just click the **And** or the **Or** button.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:92

Figure A.69 Starting to Make a Custom Filter in Etherpeek

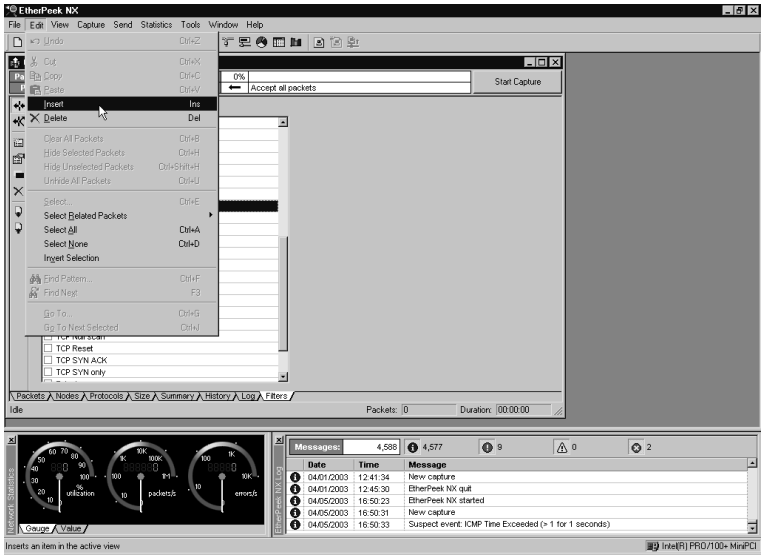

Figure A.70 Editing Etherpeek Filter Dialog Box

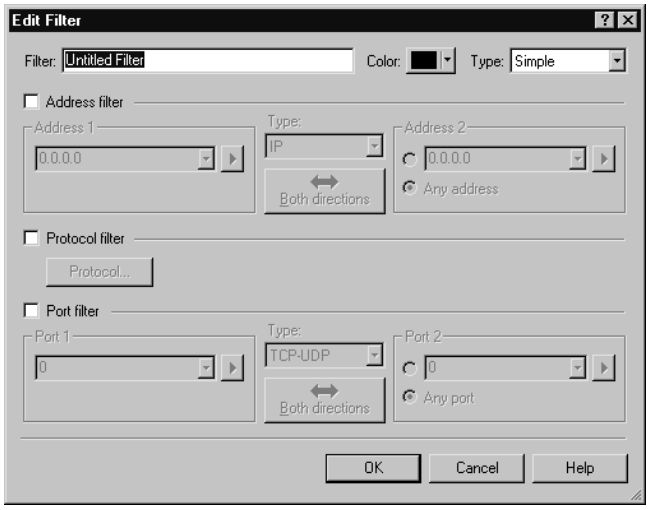

When we click the button, we are presented with a dialog box that asks on what we want to filter: ports, protocols, error, or string. Choose **string**, and then all you have to do is to type in the string and any offset that might be needed, as shown in Figure A.72.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:93

Figure A.71 Adding Boolean Logic Functions to Our Filter

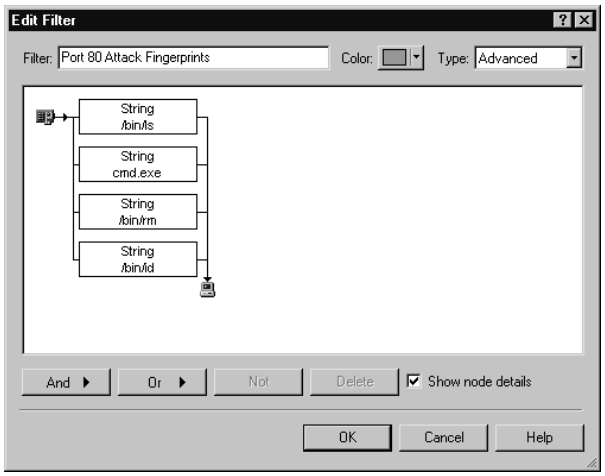

Figure A.72 Adding a String to Custom Filter

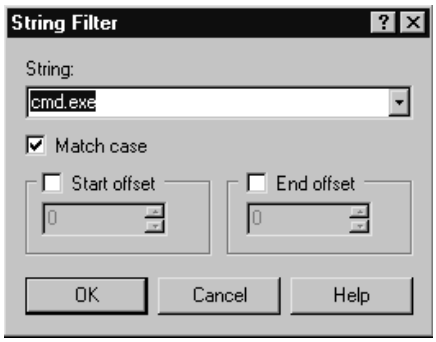

Once you have completed our filter, it will be in the list of filters to choose from. Now all you do is click the box next to our filter to enable it, as shown in Figure A.73, and start the capture. That's not too much trouble to make a custom filter for a poor man's IDS system, is it?

Network Time

A few extra points to know are that you need a plan for incident reporting and a plan for post mortem evaluations. If there will be legal action, you need to know the procedures on how to save evidence, copy log files, and report the incident. We saw previously that a way to beat the IDS and other security devices in the DMZ is to break into the DMZ slowly and then over a period time have the log files overwrite your tracks. In order to file a legal action, you need those log files, and they need to be complete and accurate. Given how

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:94

cheap disk space is nowadays, there is no excuse for not maintaining good log files. In order to aid in the tracking of the attack, a good timeline is needed. This becomes even more important when the attack might start with the DNS server in the DMZ and then jump to the Web servers or other DMZ hosts. Often the firewalls and IDS have internal clocks but they have not been set or have drifted. Correcting this normal drift of the time is where having an NTP server or atomic clock server is helpful. You can now have very accurate logs across multiple nodes on the DMZ or network using these types of technologies, and if the attack spanned hours or geographic locations, the accurate clock will be of great help in the analysis of the attack.

Figure A.73 Enabling the New Custom Filter

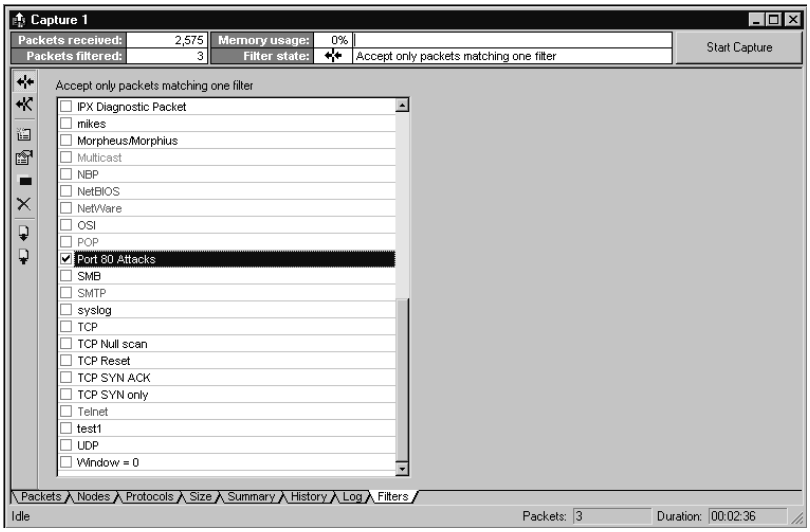

More IDS Deployment Strategies

We saw early on the classic deployment of the typical IDS sensor just in front of the firewall. We also learned that we can have a host-based IDS system. But is that all there is to deployment of an effective IDS security model? Is this the best we can do? No; this is just the start of what we can do to protect our DMZ hosts and networks. In Figure A.74 we see a mixed deployment of NIDS in the form of IDS Sensors 1 and 2 plus the mail server and the DNS server in the DMZ being protected by a HIDS system.

In this design, the IDS Sensor 1 monitors anyone who is at the Internet side of the firewall and starting to knock against it. This would be the first port scans against your network, for example. IDS Sensor 2 is positioned to watch intranet network traffic, which could be a certain segment, server farm, or the entire LAN looking for an internal attacker. As we saw earlier with the FBI report, often an attack is started on the inside of the network from someone who is trusted. In the DMZ, we have our mail server and our DNS server both

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:95

protected by a HIDS. This guards our systems against an attack that might have slipped past the firewall and IDS Sensor 1. Remember, we want layered network defense, just like the layers of an onion, and here we have a good start at it. More layers would include the ACLs on the router between the DMZ and the LAN and would even include our password policies, server permissions, and so on.

Figure A.74 Mixed NIDS and HIDS Deployment

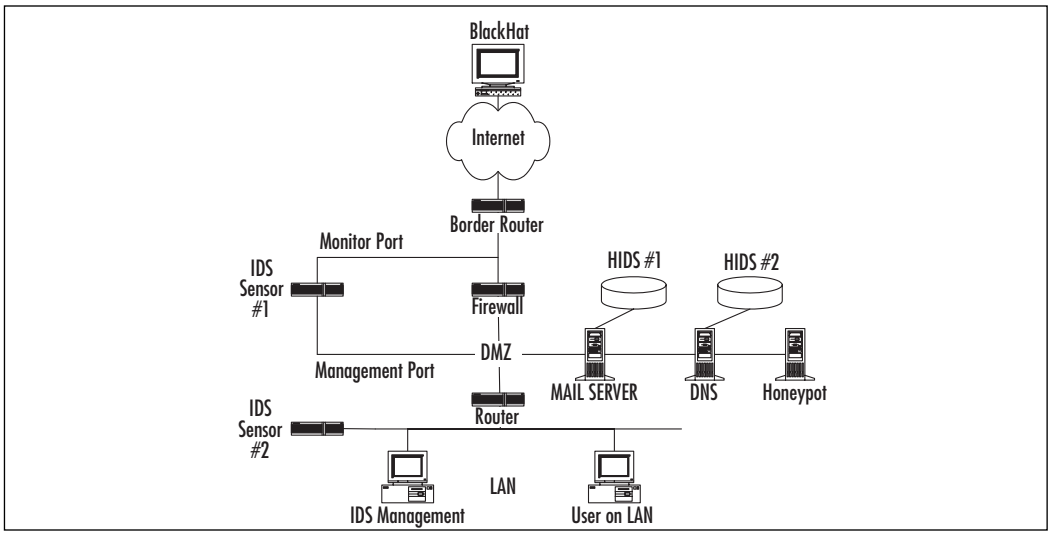

As good as this design is, we can improve it even more, as shown in Figure A.75. Remember our good friend the honeypot? Here we see a honeypot deployed in our DMZ.

Figure A.75 Mixed NIDS and HIDS with the Honeypot

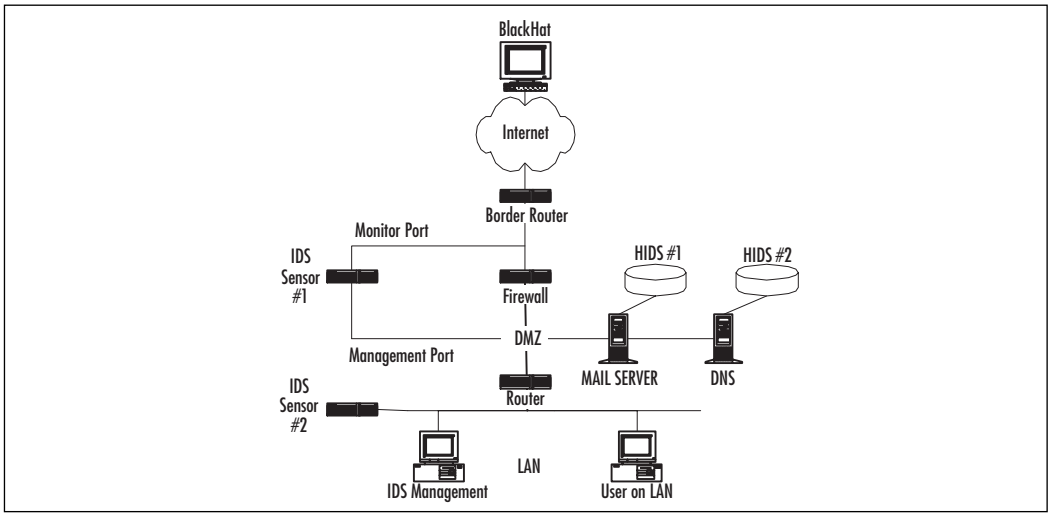

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:96

Now we really can have some fun with our network defense. We have the firewall to stop the average wannabe hacker. We have an IDS acting as the sentry, watching the front door to alert us to anyone who starts to find a way in. We have a DMZ to keep network-accessible hosts off our LAN. We have those DMZ hosts protected by a HIDS that monitors the hosts locally for any changes that might take place. We have now added the temptation of the honeypot, which we hope will keep the attacker's attention away from our real hosts and give us time to set up a defense, such as shunning his IP address, blocking his IP address, sending TCP resets, and logging his activities. On the LAN side, we have a NIDS watching for anyone who either gets past all the DMZ defenses or has started on the inside and is working his way out. We now have a layered defense configured for our DMZ and network.

Case Study

Our client, a company called Magic Potter, has a basic network installed and running. The IT staff at the company use a conventional design of a LAN with a DMZ separated by a firewall on either end of the DMZ. Within the DMZ they have their mail server, a primary and secondary DNS server, a SQL server and three Web servers. Their security officer, who is also their network administrator, Cisco geek, and desktop support technician, managed to convince her management that outside help was needed to bring up their network security standards to reflect the new threats she has been reading about in the industry trade magazines. We enter the company as security consultants. Like many IT departments, the budget is lean for this project, but after sitting down and discussing the current status of the network, future plans, support issues, and budget concerns, we assure our client that we can offer some reasonable solutions that won't break their bank too badly.

Magic Potter does not have the budget for a penetration test, but the security officer shows us the results of her own private audits using Nmap from her home PC against the company's firewalls. She has found a few holes in the firewall, but she is very concerned about vulnerabilities, such as various attacks on the IIS Web servers, and she asks us about something she has heard about, called SQL injection attacks.

We take all this down and arrive back on site a week later with some recommendations to what Magic Potter can do to help tighten things:

1. We suggested that we or Magic Potter should patch and service-pack all the servers in the DMZ to bring them up to current levels.
2. We suggested that they upgrade the DNS servers to the latest version of BIND.
3. We suggested that they split the DNS servers into one server in the DMZ and one on the internal network.
4. We suggest that they apply access lists on the firewalls to control DNS zone updates to and from certain IP addresses.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:97

5. Since the budget is lean, we suggest that they use a Snort sensor in the classic deployment, with the monitoring port watching over the firewall facing the Internet.
6. We then suggest that they use the open source Linux version of Tripwire on the Web servers, the SQL server, the mail server, and the one DNS server in the DMZ.
7. We suggest that the firewalls be configured to use only SSH for management access and that all services not required on the DMZ servers to be either removed or disabled.
8. We also suggest that they follow published best practices standards for hardening their DMZ servers. We recommend the Center for Internet Security at www.cisecurity.org as a good starting point for these standards.
9. We suggest that all the severs and firewalls send log files back to a syslog server and that Magic Potter purchase Sawmill to build reports from the syslog files.

All of these suggestions would bring Magic Potter into a much improved defensive posture at a minimal cost to the IT department. None of the suggestions required special training or expensive class time in order to be effective. We did suggest that at a later time the company might want to consider a honeypot inside the DMZ to help lure potential intruders away from the production servers.

Lessons Learned

In the last chapter of the book, Designing and Building Enterprise DMZs, and in this chapter we have covered an incredible amount of material about how to hack your DMZ, firewall, and network. Then we covered how to find the hacking, how to stop it, and how we might prevent it from happening again. The one thing that should stand out by now is that although firewalls are still a requirement on a network, the IDS is a tremendous aid to protecting your network. You need to protect the DMZ and network in layers. No single silver bullet will protect the DMZ and network.

An important lesson that should be apparent by now is how important it is to have the different parts of the network correctly architected and the equipment configured. Our mantra is, "Defaults are *not* your friend." You can have the best firewall in the world, but without defense in layers such as the DMZ, hardened routers, hardened servers, VLANs, and desktop security, in reality your defenses are very weak. Network security is constantly evolving, and to stay current on the various threats requires some effort on your part; you need to stay abreast of current events in network security. This means reading some of the newsgroups, hit the magazines, and studying what the experts such as SANS and CERT are recommending for best practices.

We have seen that once the network defenses are breached, the attacker can cause all kinds of havoc on the network—from erasing files to stealing intellectual property. Part of

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:98

the planning of a solid defense assumes the worst and that the attacker will make it past all the defenses in place and compromise something of value. You must include in your planning a business continuity plan that would be a backup strategy, server recovery plans, and how to handle the incident response, which can include legal issues such as evidence preservation. With the knowledge gleaned from these preceding chapters you can make life as difficult as possible for the would-be attacker.

Summary

As we have seen in this chapter, to prevent unwanted visitors on our network we must take more sophisticated steps than simply installing a firewall. The lesson in Chapter 14 was that hackers are a very creative group of people, and it takes someone just as creative to effectively deter them from breaking a network.

With respect to Intrusion Prevetion Systems, sensors cannot be used efficiently as standalone devices in the enterprise network. When a network and its sensors grow in size and number, the administrative overhead of the sensors becomes an ever increasing burden. When deployed in large numbers on an enterprise network, the sensors require the IPS Management Center to provide the group management functions needed for scalable operations. The IPS Management Center can group together sensors with similar configurations so that the same operations can be performed on all sensors within a group. Similarly, the MC for IPS Sensors can efficiently update the sensor software version and the signature release level of all or selected sensors in one operator action. The MC for IPS Sensors is integrated with an IPS Database where the configuration and signature settings of all of the sensors are stored. This database permits the MC for IPS Sensors operator to easily review, edit, approve, and deploy configuration settings and signature parameters for each and every sensor.

The MC contains report generation features that can be automated. Reports can be scheduled for generation at periodic intervals and can be viewed online, exported to an HTML file or posted on a company intranet. Finally the MC for IPS Sensors has various self-administration capabilities including the capability to log audit records of its own internal functions. It can even be configured to take action when certain event thresholds are reached such as the IPS database size growing beyond a configured limit.

Not only can we use applications such as Snort or combination hardware/software such as the Cisco IDS solution to monitor our network, but we can reuse other tools such as the network analyzer. With custom filters, we can turn a packet sniffer into a poor man's IDS for that quick look that we might need. IDSs such as the Cisco solution can actually monitor the IDS, routers, and PIX firewalls and then adjust the various rules and access lists on the fly to defeat attacks while they occur.

But for all the sophistication of an IDS it means nothing without good reporting. We have to read the reports generated by our security devices in order to be aware of what is happening. We also need to understand what is normal for our network and what is

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:99

abnormal. Along with the reporting comes the need to have an accurate timestamp on the logs. We can get this by using an NTP server or an atomic clock on our network. Everything should be synchronized to this timepiece so that the log files and reports all match up very nicely across the timeline.

To protect the DMZ and to have the DMZ protect the network is no small undertaking. It involves examining and securing the network, all the way from the big-picture enterprise view down to the individual desktop.

Solutions Fast Track

Intrusion Detection 101

- There are two primary types of IDS. One is network based, or NIDS, and the second is host based, or HIDS.
- HIDS provide in-depth logging, are better at detecting unknown attacks, and have fewer false positives than NIDS. However, HIDS require an agent on each host.
- The NIDS has a lower overall cost since it is an appliance instead of an agent on each host. NIDS are generally easier and quicker to deploy than HIDS.
- To help the NIDS see all traffic, you can either mirror the switch ports or use a hardware device called a *network tap*
- .■ To use IDS effectively to protect your network and the DMZ, you should consider a mix of NIDS and HIDS.

Repelling the Hacker

- One of the easiest tools to help repel a hacker is to deploy a honeypot to distract him from your real systems.
- An effective DNS solution is to split the normal two DNS servers into one DNS server in the DMZ and the second on the internal network. Use ACLs to control the flow of zone information between them.
- Use a HIDS such as Tripwire to lock down and protect your DMZ hosts.
- One of the quickest and simplest fixes in the DMZ is to keep all servers current on service packs and hotfixes.
- Follow best practices for setting up servers, such as hardening them, setting strong passwords, using SSH instead of Telnet, and other security-minded practices.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:100

- Do not give away any information to reconnaissance artists about your DMZ or hosts for free. Look into web server-masking software, use options in BIND to change version strings, and don't use "gimme" names in the DNS server, such as myWindows2K3server1.

Understanding the Cisco IPS Management Center

- The MC for IPS Sensors logs internal audit records pertinent to the Intrusion Detection System.
- The MC for IPS Sensors can manage approximately 300 sensors.
- Sensor and signature configuration are key functions performed by the MC for IPS Sensors.
- Maintaining current sensor software and signature releases are functions of the MC for IPS Sensors.

Installing the Cisco IPS Management Center

- Prerequisite products include Windows and Cisco Works Common Services.
- A related product is the Security Monitor that displays real time alarms from the sensors.

Setting up sensors and sensor groups

- Sensors should be placed at entry points to the network and between sub-networks of different security levels.
- Sensors with similar configuration settings can be placed in the same sensor group or subgroup.
- A sensor can be placed behind a filtering router so the sensor can issue a blocking command to the router when an attack is detected.

Configuring signatures and alarms

- There are six classifications of signatures: general, TCP, UDP, string-matching, ACL and custom.
- Signature settings can be configured and tuned by the MC for IPS Sensors.
- The MC for IPS Sensors can generate, approve and deploy sensor configuration files.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:101

Configuring reports

- The MC for IPS Sensors has 6 audit log reports: subsystem, sensor version import, sensor configuration import, sensor configuration deployment, console notification, and audit log.
- Reports can be generated immediately, scheduled at a later time or scheduled at regular intervals.
- The generated reports can remain online for viewing or deleted.
- The generated reports can be exported into an HTML file.
- The scheduled report parameters can be edited.

Administering the Cisco MC for IPS Sensors Server

- Database Rules are designed to trigger actions when specified database event thresholds are reached.
- The MC for IPS Sensors can be used to update sensor software versions and signature releases.
- An email server can be specified for the MC for IPS Sensors to use to distribute email notifications.

Snort

- Snort is a freely available application and extremely flexible architecture. There are many log file analyzers, graphical front ends, interfaces to SQL databases, and rule generators.
- Snort rules consist of a single line of code that can set up filters based on IP addresses, port numbers, direction of traffic, search string, or offset within the packet.

The Poor Man's IDS

- A network analyzer such as WireShark, Etherpeek, or NAI Sniffer can be used as a NIDS with custom filters.
- To build custom filters, you can find the various search strings, packet offsets, and other information from Snort rules.
- One of the easiest network analyzers to build custom filters for is Etherpeek due to the "cut and paste" function within the filter-building mechanism.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:102

More IDS Deployment Strategies

- Never rely on a single IDS for deployments to high risk networks. It might be important (based on your risk assessment) to run a combination of both HIDS and NIDS.
- An advanced strategy is based on the design of the DMZ: ACLs on the Internet router, the use of authentication, authorization, and accounting (AAA), auditing logging, a firewall, a honeypot, and last but not least, an IDS. Make sure you use NIDS where applicable, and where you can afford it or risk demands it, use HIDS on all DMZ hosts.

Lessons Learned

- Always stay current with your software patches and hotfixes to help protect your servers.
- To protect your DMZ, you need to use a layered approach, like an onion, instead of one hard layer like an egg, which is easily broken. Remember, defense-in-depth.
- You do not have to spend a lot of money in order to protect your DMZ and network. Many high-quality tools are available in open source.

Understanding the Cisco IPS Management Center

- The MC for IPS Sensors logs internal audit records pertinent to the Intrusion Detection System.
- The MC for IPS Sensors can manage approximately 300 sensors.
- Sensor and signature configuration are key functions performed by the MC for IPS Sensors.
- Maintaining current sensor software and signature releases are functions of the MC for IPS Sensors.

Installing the Cisco IPS Management Center

- Prerequisite products include Windows and Cisco Works Common Services.
- A related product is the Security Monitor that displays real time alarms from the sensors.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:103

Setting up sensors and sensor groups

- Sensors should be placed at entry points to the network and between sub-networks of different security levels.
- Sensors with similar configuration settings can be placed in the same sensor group or subgroup.
- A sensor can be placed behind a filtering router so the sensor can issue a blocking command to the router when an attack is detected.

Configuring signatures and alarms

- There are six classifications of signatures: general, TCP, UDP, string-matching, ACL and custom.
- Signature settings can be configured and tuned by the MC for IPS Sensors.
- The MC for IPS Sensors can generate, approve and deploy sensor configuration files.

Configuring reports

- The MC for IPS Sensors has 6 audit log reports: subsystem, sensor version import, sensor configuration import, sensor configuration deployment, console notification, and audit log.
- Reports can be generated immediately, scheduled at a later time or scheduled at regular intervals.
- The generated reports can remain online for viewing or deleted.
- The generated reports can be exported into an HTML file.
- The scheduled report parameters can be edited.

Administering the Cisco MC for IPS Sensors server

- Database Rules are designed to trigger actions when specified database event thresholds are reached.
- The MC for IPS Sensors can be used to update sensor software versions and signature releases.
- An email server can be specified for the MC for IPS Sensors to use to distribute email notifications.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:104

Bibliography

Barman, Scott, *Writing Information Security Policies*, (2nd Ed), New Riders, Indianapolis, IN., 2002

Pfleeger, Charles P., *Security in Computing*, (2nd Ed), Prentice Hall PTR, Upper Saddle River, NJ., 1997

SANS – Security Policy Project, http://www.sans.org/resources/policies/

NIST – "Guidelines on Firewalls and Firewall Policy", NIST, http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf

National State Auditors Association and U.S. General Accounting Office – "Management Planning Guide for Information Systems Security Auditing", http://www.gao.gov/special.pubs/mgmtpln.pdf

Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to **www.syngress.com/solutions** and click on the **"Ask the Author"** form.

Q: Where in my enterprise network should I deploy my sensors?

A: In your enterprise network, the sensors should be deployed at entry points to your network and between sub-networks that require different levels of protection. This is not just internet connections but any connection to a vendors network by VPN or other connection

Q: How can I make sure that my sensors have signatures for the latest threats?

A: To be informed of the latest update files by email, you can subscribe to the Cisco IPS Active Update Notification.

Q: Will the MC for IPS Sensors protect my network from Denial of Service (DoS) attacks?

A: The MC for IPS Sensors itself will not protect your network from DoS attacks. However, you can use the MC to configure your IPS sensors to warn you of an attack and allow you to take appropriate action to filter the attack packets. Second, the IPS can configure the sensors to order the blocking router to block the attack.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:105

Q: Can the MC for IPS Sensors manage honey pots?

A: Through the use of a sensor on the honey pot network, the MC for IPS Sensors can detect attacks directed against honey pots and notify you so you can take appropriate action to determine the source and nature of the action.

Q: Does the IPS Management Center display real time intrusion alarms?

A: No. The Security Monitor will display real time intrusion alarms through its Event Viewer. SecMon will send email notification when certain event thresholds are reached.

Q: Using MC for IPS Sensors, can I update the configurations on several sensors at one time?

A: Yes. If all the sensors in the same group require the same configuration update, the configuration updates can be deployed with the same operator action.

Q: Can the MC for IPS Sensors manage sensors outside of my firewall?

A: Yes. The MC for IPS Sensors only requires a TCP connection to the sensor it manages. It is not even necessary that the sensor be in the same network as the MC for IPS Sensors.

Q: Can a large network be managed by multiple MC for IPS Sensorss?

A: Yes. Different portions of a large network may have different security policies and it may be more advantageous from an administrative perspective to manage it with more than one MC for IPS Sensors.

Q: How can I minimize false alarms?

A: By tracking and analyzing false positives that are generated, you can determine the optimal settings for your signatures to minimize false alarms. This is usually done by tuning the Micro-Engine Parameters in your signatures or by excluding certain internal networks as triggers of alarms.

Q: How can I keep my IDS from being "seen" on the network?

A: An easy way is to use an Ethernet cable with the transmit pairs cut so that the IDS only receives traffic. An optional but better method is use a network tap, which keeps even the MAC address of the IDS sensor from appearing on the network.

Q: Can I use more than one IDS sensor on a network?

A: Sure. In fact, it is highly recommended to use more than a single sensor to provide adequate protection.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:106

Q: Where in my enterprise network should I deploy my sensors?

A: In your enterprise network, the sensors should be deployed at entry points to your network and between sub-networks that require different levels of protection. This is not just internet connections but any connection to a vendors network by VPN or other connection

Q: How can I make sure that my sensors have signatures for the latest threats?

A: To be informed of the latest update files by email, you can subscribe to the Cisco IPS Active Update Notification.

Q: Will the MC for IPS Sensors protect my network from Denial of Service (DoS) attacks?

A: The MC for IPS Sensors itself will not protect your network from DoS attacks. However, you can use the MC to configure your IPS sensors to warn you of an attack and allow you to take appropriate action to filter the attack packets. Second, the IPS can configure the sensors to order the blocking router to block the attack.

Q: Can the MC for IPS Sensors manage honey pots?

A: Through the use of a sensor on the honey pot network, the MC for IPS Sensors can detect attacks directed against honey pots and notify you so you can take appropriate action to determine the source and nature of the action.

Q: Does the IPS Management Center display real time intrusion alarms?

A: No. The Security Monitor will display real time intrusion alarms through its Event Viewer. SecMon will send email notification when certain event thresholds are reached.

Q: Using MC for IPS Sensors, can I update the configurations on several sensors at one time?

A: Yes. If all the sensors in the same group require the same configuration update, the configuration updates can be deployed with the same operator action.

Q: Can the MC for IPS Sensors manage sensors outside of my firewall?

A: Yes. The MC for IPS Sensors only requires a TCP connection to the sensor it manages. It is not even necessary that the sensor be in the same network as the MC for IPS Sensors.

Q: Can a large network be managed by multiple MC for IPS Sensorss?

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:107

A: Yes. Different portions of a large network may have different security policies and it may be more advantageous from an administrative perspective to manage it with more than one MC for IPS Sensors.

Q: How can I minimize false alarms?

A: By tracking and analyzing false positives that are generated, you can determine the optimal settings for your signatures to minimize false alarms. This is usually done by tuning the Micro-Engine Parameters in your signatures or by excluding certain internal networks as triggers of alarms.

Q: Where can I find strings to build custom filters for my sniffer or rules?

A: Snort.org has one of the best databases for rules. You can look at the rule and find out exactly the string you need to filter or even the hex code and offset.

Q: I can only afford a single IDS. Where should I put it in my network?

A: The traditional deployment of a single IDS is to place the monitoring port in the DMZ and the management port in the trusted network. This way you can watch for anything that gets past the firewall's blocking. There is some debate as to whether this is optimal or the monitoring port should be on the Internet side of the firewall. Try them both and see what works best for your network.

403_Ent_DMZ_AA.qxd 10/25/06 11:23 AM Page A:108